SENSITIVE SEARCH FOR RADIO VARIABLES AND TRANSIENTS IN THE EXTENDED CHANDRA DEEP-FIELD SOUTH

K. P. Mooley[1], D. A. Frail[2], E. O. Ofek[3], N. A. Miller[4], S. R. Kulkarni[1] & A. Horesh[1]

Manuscript version March 27, 2013. Accepted in ApJ.

ABSTRACT

We report on an analysis of the Extended Chandra Deep Field South (E-CDFS) region using archival data from the Very Large Array, with the goal of studying radio variability and transients at the sub-mJy level. The 49 epochs of E-CDFS observations at 1.4 GHz sample timescales from one day to 3 months. We find that only a fraction (1%) of unresolved radio sources above 40 μJy are variable at the 4σ level. There is no evidence that the fractional variability changes along with the known transition of radio source populations below one milliJansky. Optical identifications of the sources show that the variable radio emission is associated with the central regions of an active galactic nucleus or a star-forming galaxy. After a detailed comparison of the efficacy of various source-finding algorithms, we use the best to carry out a transient search. No transients were found. This implies that the areal density of transients with peak flux density greater than 0.21 mJy is less than 0.37 deg^{-2} (at a confidence level of 95%). This result is approximately an order of magnitude below the transient rate measured at 5 GHz by Bower *et al.* (2007) but it is consistent with more recent upper limits from Frail *et al.* (2012). Our findings suggest that the radio sky at 1.4 GHz is relatively quiet. For multi-wavelength transient searches, such as the electromagnetic counterparts to gravitational waves, this frequency may be optimal for reducing the high background of false positives.

Subject headings: catalogs — radio continuum: galaxies — surveys

1. INTRODUCTION

For more than four decades the largest science yield of variable and transient radio emission has come from single-dish radio telescopes, which have surveyed the sky for pulsed and periodic emission from pulsars and related phenomena on typical timescales of seconds to milliseconds. In contrast, interferometric imaging surveys, which are best-suited to probe timescales ranging from seconds, hours and days, have only just begun. There are a limited number of surveys of these so-called *long duration* transients and variables at frequencies of 1–10 GHz, each with a different sensitivity, cadence and field-of-view (Ofek et al. 2011; Bell et al. 2011). Fully exploring this phase space is one of the main science drivers for a new generation of synoptic radio imaging facilities, such as ASKAP (Johnston et al. 2008) and Apertif/WSRT (Oosterloo et al. 2010), that are being built in the coming years (see also Lazio et al. 2009).

Our knowledge of the variable GHz sky is especially lacking at sub-milliJansky flux density levels. For *persistent* sources there is a well-known flattening of the Euclidean-normalized radio source counts below about 1 mJy, corresponding to a change in the radio source populations. This flattening is likely due to the emerging importance of star forming galaxies and low-luminosity active galactic nuclei (AGNs) at redshifts of order unity (Condon 2007; Seymour et al. 2008; Smolčić et al. 2008; Padovani 2011; Condon et al. 2012). Thus, while the

variability studies above milliJansky levels are dominated by radio-loud AGNs with compact, flat-spectrum components (Sadler et al. 2006; Lovell et al. 2008), variability at these deeper flux density levels may probe new source populations.

There are indications that the mJy transient sky is exciting. Over the last decade astronomers have detected transient decimetric emission from a variety of sources: transient, bursting and pulsed radio emission from magnetars (Cameron et al. 2005; Gaensler et al. 2005; Camilo et al. 2006), short-lived radio afterglows of short-duration gamma-ray bursts (Fox et al. 2005), emission from (a transient) jet in a dwarf nova (Körding et al. 2008), a new population of sub-relativistic supernovae (Soderberg et al. 2010), a mysterious population of bursting radio sources (Hyman et al. 2005, 2009), and relativistic outflows from tidal disruption events (Zauderer et al. 2011; Cenko et al. 2012). However, with few exceptions, most of what we know about the transient radio sky has come via radio follow-up of objects identified by synoptic telescopes at optical, X-ray or gamma-ray wavelengths. Clearly, this titillating trove of (serendipitous) discoveries call for systematic exploration of the decimetric sky on timescales of minutes and longer.

A ready source of archival data for searching for sub-milliJansky transients and variables comes from deep continuum imaging surveys of the GHz radio sky undertaken to study the extragalactic radio source populations. Many such surveys have been carried out, reaching noise levels of 4–10 μJy and with arcsecond resolution (e.g. Schinnerer et al. 2007; de Zotti et al. 2010). In order to reach these deep flux density limits, it is standard to observe these fields for many epochs with a cadence that samples timescales of days, weeks and months. An added benefit is that these deep surveys are accompa-

[1] Cahill Center for Astronomy, MC 249-17, California Institute of Technology, Pasadena, CA 91125, USA
[2] National Radio Astronomy Observatory, P.O. Box O, Socorro, NM 87801
[3] Benoziyo Center for Astrophysics, Faculty of Physics, The Weizmann Institute for Science, Rehovot 76100, Israel
[4] Department of Astronomy, University of Maryland, College Park, MD, 20742-2421, USA

nied by rich multi-wavelength continuum datasets and optical/infrared spectroscopic measurements. Thus the counterpart of any unusual variable or transient source can be readily identified and its redshift determined.

In this paper we present a search for transients and variables at sub-milliJansky flux density levels using data taken as part of a deep radio continuum survey toward a region known as the Extended Chandra Deep Field-South (E-CDFS; Miller et al. 2008). In §2 we describe the original survey and our re-reduction of the data. In §3 we describe how the variability light curves for 599 point sources were extracted. The transient search is described in §4. The interpretation and implication of these results for radio source variability and transients is discussed in §5.

2. OBSERVATIONS AND DATA REDUCTION

The E-CDFS is an intensely studied region having a plethora of multi-wavelength data available (viz. X-ray, ultraviolet, optical, infrared and radio; see Miller et al. 2008, and references therein). Here, we use data from the radio survey undertaken by Miller et al. (2008)[5]. The observations were made at a frequency of 1.4 GHz, using the National Radio Astronomy Observatory (NRAO) VLA in its A configuration (Project code AM 889). A hexagonal grid of six pointings was made, with each pointing separated 12′ from its nearest neighbor (see Table 1 and Figure 1). All the observations were carried out in 2007 between June 15 and September 23. Only a single pointing was observed for each epoch. Each epoch was a 5-hr track centered on 03:30 LST[6]. There were a total of 49 epochs, with a combined allocation of 245 hrs (see Table 2). The total number of epochs for each pointing N_{ep} is given in Table 1. For more details about the specifics of the observational setup see Miller et al. (2008).

The original purpose of these data was to average together all pointings and epochs in order to create a deep ($\sigma_{rms} = 5$–$8 \ \mu$Jy) continuum image of the E-CDFS. In order to explore variability and to search for transients, we needed to work with the single epoch, single pointing images instead. Despite this, many of the data reduction steps that we followed were similar to Miller et al. (2008). We summarize the process here, pointing out slight differences. All calibration and imaging was carried out in the Astronomical Image Processing System (*AIPS*) package[7].

The calibrated *uv* data for each epoch were imaged and deconvolved separately. In order to image the entire field-of-view at full resolution, we created 37 different facets in a flys-eye pattern. Each of the 37 facets was an image of 1024^2 pixels with a 0.5″ pixel size offset from one another. There were another 23-25 outlier fields 128^2 pixels each made of cataloged bright radio sources outside the primary beam of each pointing but within a 2° radius.

The *AIPS* task IMAGR was used to deconvolve each image to the rms noise level. In order to have a consistent

[5] In the past, radio observations of this field have also been carried out by Kellerman et al. (2008), Afonso et al. (2006), and Norris et al. (2006).

[6] This is of great benefit to variability studies, which are otherwise plagued by changes in the observing setup.
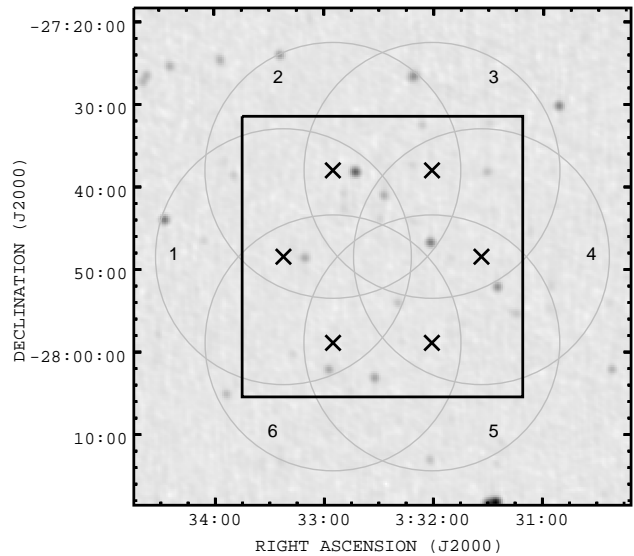
[7] http://www.aips.nrao.edu/

FIG. 1.— Hexagonal grid of six pointings (crosses), each pointing separated 12′ from its nearest neighbor. The $34' \times 34'$ (black square) region represents the extent of the final image from the Miller et al. (2008) data release. 15′-circles (grey) corresponding to the 50% beam attenuation in different pointings are also shown. The $60' \times 60'$ background image is from NVSS.

TABLE 1
LIST OF SURVEY POINTINGS

Pointing ID	R.A. (J2000.0)	Dec. (J2000.0)	N_{ep}
ECDFS1	03 33 22.25	−27 48 30.0	7
ECDFS2	03 32 55.12	−27 38 03.0	9
ECDFS3	03 32 00.88	−27 38 03.0	8
ECDFS4	03 31 33.75	−27 48 30.0	8
ECDFS5	03 32 00.88	−27 58 57.0	9
ECDFS6	03 32 55.12	−27 58 57.0	8

NOTE. — N_{ep} is the number of epochs per pointing.

set of images, we applied a Gaussian taper to the (30% level) of 100 kλ and 70 kλ in the u and v directions to the visibility data, and we restored the final images to a synthesized beam of $2.8'' \times 1.6''$ (position angle$\simeq 0°$). After deconvolution, the 37 facets for each epoch were combined using the *AIPS* task FLATN to form a single 5120^2 pixel image 42.7′ across. A correction for the attenuation from the primary beam was not applied at this stage in order that the images used in analysis had uniform noise statistics across the entire image. The rms noise σ_{rms} for each epoch is given in Table 2.

3. VARIABILITY ANALYSIS

The source catalog we used to investigate variability was taken from the second data release (DR2; Miller et al. 2013) of Miller et al. (2008). This catalog was generated by combining all the data from Tables 1 and 2 to make a single deep $34' \times 34'$ image with a typical sensitivity of 7.4 μJy. Miller et al. (2013) identified sources using the *AIPS* task SAD down to 4σ, and then inspected the residual map to identify missed sources as well as accepted sources which were poorly fit by SAD. These missing sources were then added to the preliminary source list. Further flagging and follow-up was done in order to produce a modified source list in which all sources with

TABLE 2
OBSERVING EPOCHS

Epoch	Date UT	Pointing	σ_{rms} μJy bm^{-1}
1	2007 Jun. 15	ECDFS2	26.3
2	2007 Jun. 24	ECDFS3	28.0
3	2007 Jun. 25	ECDFS4	28.6
4	2007 Jul. 01	ECDFS6	25.9
5	2007 Jul. 06	ECDFS5	29.2
6	2007 Jul. 12	ECDFS1	26.6
7	2007 Jul. 13	ECDFS2	26.6
8	2007 Jul. 14	ECDFS3	25.5
9	2007 Jul. 15	ECDFS4	26.4
10	2007 Jul. 16	ECDFS5	26.8
11	2007 Jul. 17	ECDFS6	26.1
12	2007 Jul. 19	ECDFS1	26.0
13	2007 Jul. 20	ECDFS2	34.2
14	2007 Jul. 21	ECDFS3	26.0
15	2007 Jul. 22	ECDFS4	25.3
16	2007 Jul. 23	ECDFS5	27.8
17	2007 Jul. 24	ECDFS6	27.5
18	2007 Jul. 26	ECDFS5	31.8
19	2007 Jul. 27	ECDFS2	27.2
20	2007 Jul. 28	ECDFS3	30.5
21	2007 Jul. 29	ECDFS4	27.1
22	2007 Jul. 30	ECDFS1	28.0
23	2007 Aug. 02	ECDFS6	27.0
24	2007 Aug. 03	ECDFS1	27.4
25	2007 Aug. 04	ECDFS2	26.7
26	2007 Aug. 05	ECDFS3	25.9
27	2007 Aug. 06	ECDFS4	29.1
28	2007 Aug. 09	ECDFS5	31.2
29	2007 Aug. 10	ECDFS6	30.0
30	2007 Aug. 11	ECDFS1	31.9
31	2007 Aug. 13	ECDFS2	31.3
32	2007 Aug. 14	ECDFS3	31.5
33	2007 Aug. 16	ECDFS4	30.8
34	2007 Aug. 17	ECDFS5	45.7
35	2007 Aug. 18	ECDFS6	28.6
36	2007 Aug. 21	ECDFS1	29.1
37	2007 Aug. 23	ECDFS2	30.9
38	2007 Aug. 25	ECDFS3	30.8
39	2007 Aug. 26	ECDFS4	29.8
40	2007 Aug. 28	ECDFS5	33.9
41	2007 Aug. 31	ECDFS6	29.3
42	2007 Sept. 06	ECDFS2	30.9
43	2007 Sept. 07	ECDFS3	29.3
44	2007 Sept. 08	ECDFS4	29.3
45	2007 Sept. 09	ECDFS5	28.9
46	2007 Sept. 10	ECDFS6	30.2
47	2007 Sept. 11	ECDFS1	30.4
48	2007 Sept. 12	ECDFS5	37.9
49	2007 Sept. 23	ECDFS2	31.6

NOTE. — List of the 49 epochs. Each epoch consisted of a 5-hr track centered on 03:30 LST.

peak flux density greater than 5 times the local rms noise (i.e. 5σ), were fit using the *AIPS* task JMFIT. Also, the effect of bandwidth smearing on sources within the six individual pointings was assessed using JMFIT, and the resolution information was thus preserved in the output catalog. Lastly, the sources in the DR2 catalog were compared with the catalog of Kellerman et al. (2008).

The DR2 catalog contains almost twice as many sources (883 vs. 464) compared to the first data release (see Miller et al. 2013, 2008) owing to a more careful data reduction. Of the 883 sources in the DR2 catalog we created a point-source-only catalog of 736 objects used for exploring variability. With this careful approach outlined above, we expect the Miller et al. DR2 catalog to contain all real sources above 5σ (however, see §4.1).

We thus justify our use of the DR2 catalog for investi-

gating the variability of the sub-milliJansky population. In §4 we use both the DR2 image and its source catalog as a testbed for different source-finding algorithms. The signal-to-noise ratio (SNR) of sources in the DR2 catalog is shown in Figure 2.
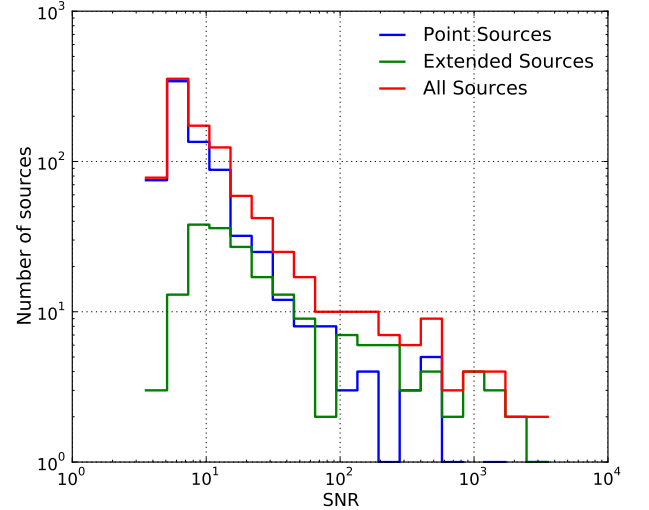

FIG. 2.— A histogram of the signal-to-noise ratio (SNR) of the sources in the DR2 catalog.

It is important to quantify the uncertainty in the peak flux density when analyzing variability of sources in different epochs. Following Thyagarajan et al. (2011), we attribute this uncertainty primarily to a combination of four[8] causes: (i) local image rms noise, (ii) uncertainty in the primary beam, (iii) flux density calibration amplitude, (iv) pointing errors. Let us denote the measured peak flux density of a source in epoch i by f_i, such that the primary beam correction is $b = b(\theta)$, where θ is the angular distance of the source from the phase center. We wish to calculate the uncertainty in the quantity (f_i/b). The image local noise (Δf_i) scales as $\Delta f_i/b$. If we denote the fractional uncertainty in the beam as ϵ_b and that in the flux density calibration amplitude by ϵ_c, then the corresponding errors scales as $(f_i/b)\epsilon$. We adopt a value of 4% for ϵ_c, intermediate between the conservative estimate of Thyagarajan et al., 5%, and the one quoted by Ofek et al. (2011), 3%. The typical pointing error $(\Delta\theta)$ of a VLA antenna is between $10''$–$20''$. The resultant uncertainty scales as $(f_i/b^2)(-db/d\theta)\Delta\theta$. All these four error terms, added in quadrature, would give the total uncertainty, σ_i, in the peak flux density corrected for the primary beam attenuation. However, as shown below, the pointing-related error term is much smaller than the rest, and hence can be neglected. Thus,

$$\sigma_i = \frac{1}{b}\sqrt{\Delta f_i^2 + f_i^2(\epsilon_b^2 + \epsilon_c^2)} \qquad (1)$$

Polynomial coefficients (and the associated error) that

[8] Far out in the primary beam, orthogonally polarized beams can be far offset on the sky, making amplitude calibration difficult. For short observations, the effects of this "beam squint" can be much larger than all other uncertainties combined. However, these offsets tend to average out when observations are made over sufficiently long times. Since the VLA's beam squint is oriented almost exactly East-West (Cotton & Perley 2010), our observational setup is optimal for averaging out the effects of beam squint.

express the average angular dependence of b can be found in the AIPS task `PBCOR`, while measurements of the VLA beam power response to beyond the first null are given in Cotton & Perley (2010). Using the beam response profile from Cotton & Perley , we can estimate the error terms (i)–(iv) above for a typical source in the DR2 catalog, having measured flux density of $300 \pm 30 \ \mu$Jy. If the source lies at the half-power radius ($\theta = 15'$), then these correction factors are about 10%, 4%, 4% and 1% of the primary-beam-corrected flux density ($570 \ \mu$Jy) respectively. The pointing-related uncertainty is thus negligible.

In light of Ofek et al. (2011), we use two measures of variability (see also Scheers 2011), the modulation index defined as the standard deviation divided by the mean,

$$m = \frac{1}{\bar{f}} \sqrt{\frac{1}{N-1} \sum_{i=1}^{N} (f_i - \bar{f})^2}, \qquad (2)$$

and the χ^2,

$$\chi^2 = \sum_{i=1}^{N} \frac{(f_i - \bar{f})^2}{\sigma_i^2}, \qquad (3)$$

where N is the number of epochs, and $\bar{f}$ is the mean flux density of the source over all the epochs considered for the variability analysis. In these two equations, the primary beam correction is implicit in f_i and $\bar{f}$. χ^2 gives a measure of the deviation from stochastic epoch-to-epoch fluctuations in the peak flux density, and we define "significant variability" beyond a level[9] of 4σ. The modulation index indicates the strength of variability, i.e. the fractional variation of the peak flux density.

3.1. Single-Pointing Variability

A variability analysis was carried out on each pointing in Table 1 separately. Peak flux densities were measured for all point sources in the DR2 catalog brighter than 40 μJy and within a $15'$ radius of the pointing centers (i.e. the 50% response radius of the primary beam of the VLA antennas; see also Figure 1). This approach has the merit of being simple and robust. Since the angular distance of a source from its pointing center is constant, the accuracy of the correction for the primary beam attenuation $b(\theta)$ is unimportant. The modulation index and χ^2 measures of variability are insensitive to a constant $b(\theta)$.

The limitation of this approach is that the resulting light curves are constructed for only 7–9 epochs. Many of our investigated point sources are found in multiple pointings and therefore light curves can be constructed with many more epochs, resulting in higher cadence over the full 100 days of observing. A full variability analysis of this kind is carried out in §3.2.

In Figure 3 we show variability plots (i.e. χ^2 vs. m) for the DR2 sources in each of the six pointings. There are approximately 175 sources per pointing. Some of the bright (> 3 mJy) sources show significant variability but with only low modulation indices ($\sim$10%). We

---

[9] For Gaussian noise, 4σ corresponds to a probability of about 1/16,000, while the number of measurements in our variability analysis (several to tens of epochs multiplied by a few hundred sources) ranges from about 1,500 to 15,000.

define strong variables as sources having $m > 0.5$ (i.e. higher than 50% fractional variability). Only two genuine variables were found in this single-pointing analysis; no strong variables were found (Table 3, upper panel).

3.2. Full Variability Analysis

In order to undertake variability analysis using data from all the epochs, the single-epoch peak flux densities were required to be corrected for (i) bandwidth-smearing, and (ii) the primary beam response, as any given source will lie at different angular distance from the phase center in different pointings. We applied bandwidth smearing correction from the approximation given in Bridle & Schwab (1989) (Eqn. 13-19 therein).

There were two choices for the primary beam profile — one derived as the empirical beam profile for VLA-FIRST (Thyagarajan et al. 2011), and the other found in the AIPS task `PBCOR`. To test which of these profiles best represented our data, we adopted the following approach. We first normalized the peak flux densities of $> 300 \ \mu$Jy sources from all epochs using their peak flux density from the DR2 catalog and plotted them as a function of distance from the pointing center. The resultant beam profile matched with the VLA-FIRST profile better than the one from `PBCOR` (to within 1%, but only for $\theta < 12'$; scatter of 6.5%). Hence we used the former beam profile for our primary beam correction, $b(\theta)$; the associated error (ϵ_b), was also taken from Thyagarajan et al. (2011). Thus, for a reliable all-epoch variability analysis, we restricted our search to the point sources in the DR2 catalog which were located within $12'$ from the pointing centers of their respective epochs. This also appears to be the radius beyond which our bandwidth- smearing approximation starts to break down. Thus, for example, a source located at $\alpha = 03^{\mathrm{h}}33^{\mathrm{m}}00^{\mathrm{s}}$ and $\delta = -28°00'00''$ would be present in pointings 6, 5, and 1, but not in 2,3, and 4 because the separation between the source and the centers of pointings 2,3, and 4 is larger than 12 arcmin. Further, as we did with the single-pointing variability (§3.1), we restricted our analysis to sources whose mean flux density was brighter than 40 μJy. This full variability analysis was carried out on 599 point sources.

Depending on the number of pointings in which a source is present, this analysis allowed us to exploit the higher cadence over the entire duration of the observing program. The resulting light curves are now more densely sampled with 15–26 epochs, rather than the 7–9 epochs for the single-pointing variability case. In Figure 4 we show variability plots for the DR2 sources, taking into account all the epochs.

Seven significant variables were found via this procedure, but no strong variables ($m > 0.5$). Both of the variables identified in the single-pointing analysis in §3.1 are also seen here. The results of the variability study are given in Table 3 and the light-curves for the significant variables are shown in Figure 5. We can compare our variability criteria with the Carilli et al. (2003) measure for variability, i.e. $\mathrm{V}_c = (S_1 - S_2)/S$, where S_1 and S_2 are the maximum and minimum flux densities respectively observed, and S is their mean. This measure of variability for the seven variable sources found in this work are listed in Table 3.

In order to undertake multiwavelength identifications we had to align the reference frames of all the data sets.
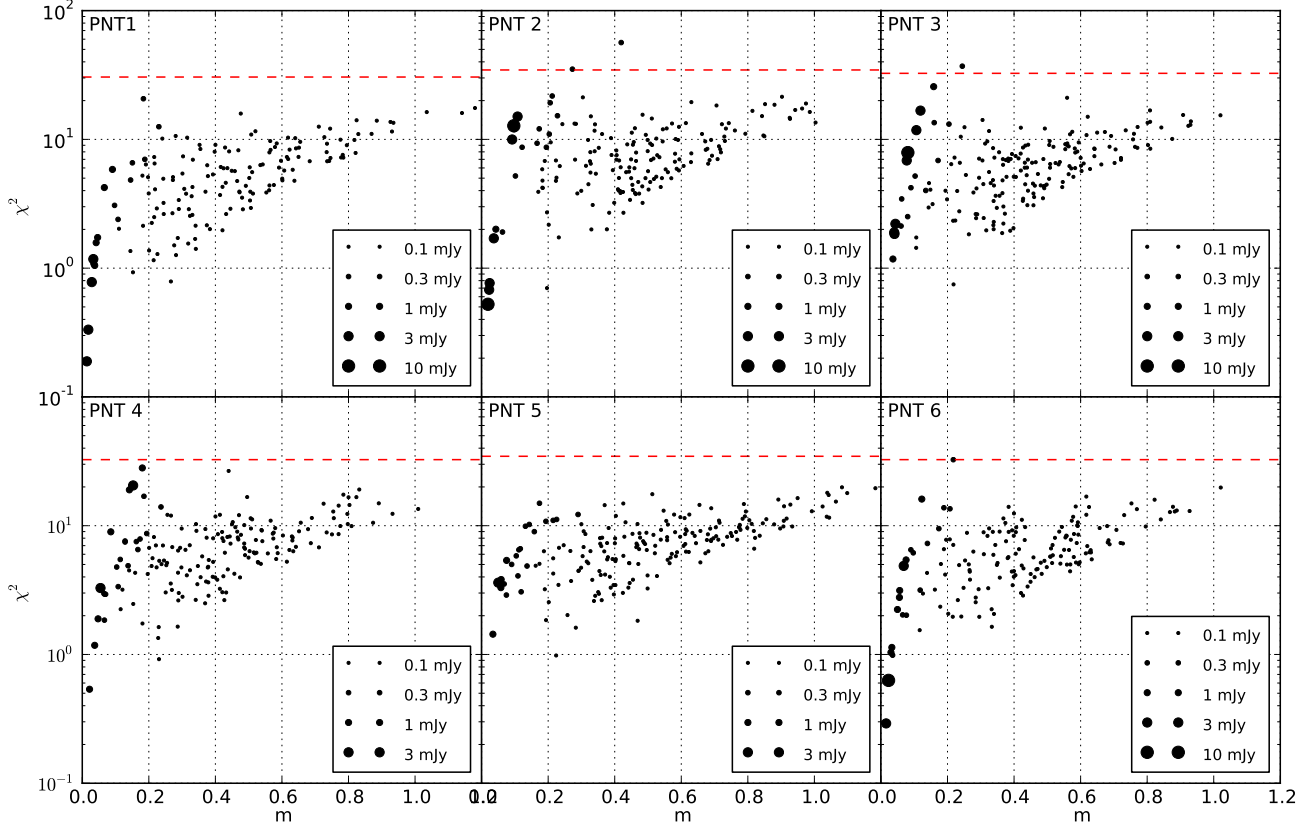
FIG. 3.— Variability plot for the point sources in the Miller et al. DR2 catalog, shown separately for all pointings. The peak flux density is denoted by the symbol size. The red dashed line represents the 4σ level for the appropriate number of degrees of freedom (one less the number of epochs in each pointing) for each pointing. The number of epochs in each pointing lies between 7 and 9.

TABLE 3
VARIABLES AMONG MILLER ET AL. DR2 SOURCES

ID	α_{J2000} (h, m, s)	δ_{J2000} (°, ′, ″)	$\bar{f}$ (μJy)	m	V_c (%)	N_{ep}	z	$\log \bar{L}_R$ (cgs)	α_R	α_{IR}	M_R (mag)	Energy Source
					Single Pointing Analysis							
337	03 32 18.03	-27 47 18.8	558 ± 14	0.26	76	8						
621	03 33 15.00	-27 51 51.3	663 ± 22	0.23	93	9						
					Full Variability Analysis							
67	03 31 27.07	-27 44 09.9	136 ± 9	0.50	181	16	1.005	30.6	$+0.23$	> -3.8	-21.04	SF+AGN
190	03 31 52.13	-27 39 26.6	891 ± 12	0.15	48	16	2.296	32.1	-0.35	> -1.3	-22.07	AGN
239	03 32 00.85	-27 35 57.1	1822 ± 27	0.12	38	8	0.266	30.5	-0.30	> -2.3	-17.26	SF+AGN
297	03 32 11.66	-27 37 26.3	3600 ± 36	0.11	41	17	0.605	31.5	$+0.89$	-0.9	-23.98	AGN
337	03 32 18.03	-27 47 18.8	494 ± 11	0.26	95	16	0.734	30.8	-0.08	> -3.4	-22.69	SF+AGN
621	03 33 15.00	-27 51 51.3	497 ± 10	0.20	70	15	1.107	31.2	$+0.24$	> -3.4	-20.71	SF+AGN
628	03 33 16.74	-27 56 30.4	1341 ± 16	0.11	38	15	0.685	31.2	-0.40	-2.2	-21.21	SF+AGN

Notes— (1) ID is the source-ID as given in the Miller et al. DR2 catalog. (2) $\bar{f}$ refers to the mean flux density corrected for the primary beam and bandwidth smearing. (3) V_c is the Carilli et al. (2003) variability criteria as described in § 3.2. (4) Redshift z is according to Bonzini et al. (2012) or Treister et al. (2009). The redshifts of ID 67, ID 239 and ID 297 as per the COMBO-17 survey catalog (Wolf et al. 2004) are 0.548, 0.947 and 1.574 respectively. (4) The 1.4 GHz spectral luminosity (erg s^{-1} Hz^{-1}) is $\bar{L}_R = 4\pi d_l^2 \bar{f}/(1+z)$, where the luminosity distance d_l assumes cosmological parameters from Komatsu et al. (2011). (5) The spectral indices ($S \propto \nu^\alpha$) between 1.4 GHz and 5.5 GHz and between 24 μm and 70 μm are tabulated as α_R and α_{IR} respectively. In the absence of a 70 μm counterpart, a 3σ upper limit to the flux density at this wavelength is considered. Note that the 1.4 GHz and 5.5 GHz measurements are non-simultaneous. (6) The absolute R-band magnitude M_R has been calculated using the redshift z and the apparent magnitude from the NASA/IPAC Extragalactic Database. (7) The energy sources within the galaxy as indicated by the radio and mid-to-far-infrared properties are listed in the last column. SF: star-formation, AGN: active galactic nucleus (see §3.3).
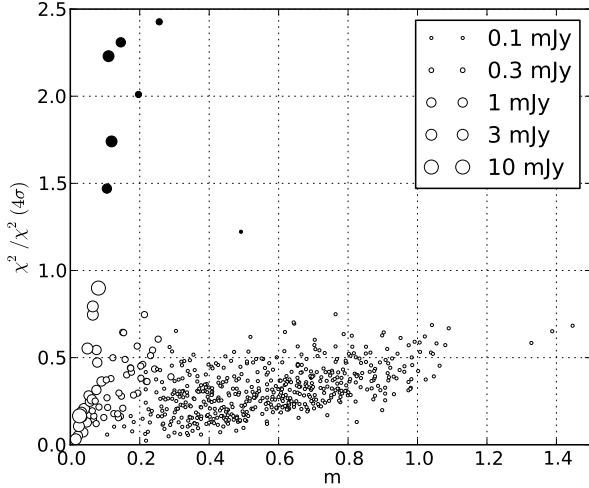
FIG. 4.— χ^2 normalized by its value at the 4σ level, plotted against the modulation index, m, for sources in the Miller et al. DR2 catalog using peak fluxes from all pointings, and corrected with empirically-derived beam attenuation profile. The 4σ level is different for different sources, depending on the number of epochs in which they are present. The mean peak flux density is denoted by the symbol size. Filled circles indicate significant variables (lying above a normalized χ^2 of unity).

The radio and optical source positions were brought to the same reference frame by calculating the radio-source position offsets with respect to *Hubble* Space Telescope source positions from GOODS-S (Giavalisco et al. 2004) and the GEMS (Rix et al. 2004). Optical counterparts (from these two HST catalogs) were searched toward radio sources within $1''$. Only point-like sources having a single counterpart were chosen. A histogram of the offsets of these counterparts in right ascension ($\Delta\alpha = \alpha_{\mathrm{radio}} - \alpha_{\mathrm{optical}}$) and declination ($\Delta\delta = \delta_{\mathrm{radio}} - \delta_{\mathrm{optical}}$) was then computed to find the most-likely offset (peak of the histogram). The associated error was taken to be the standard deviation about this most-likely offset added in quadrature with the cataloged mean positional error of the radio source. In addition to the Miller et al. (2013) DR2 catalog, we repeated this procedure for other radio-source catalogs in the E-CDFS region, viz. Kellerman et al. (2008), Norris et al. (2006), and Afonso et al. (2006) using $\geqslant 5\sigma$ sources. The computed radio vs. optical positional offsets along with the mean positional errors listed in the respective radio catalogs are shown in Table 4.

TABLE 4
RADIO POSITIONAL OFFSETS WRT. OPTICAL HST CATALOGS

	arcsec	(1)	(2)	(3)	(4)
	σ_α	~ 0.1	0.37	0.31	[b]
	σ_δ	~ 0.1	0.57	0.58	[b]
GEMS	$\Delta\alpha$	0.18 ± 0.31	0.15 ± 0.45	-0.09 ± 0.41	-0.17 ± 0.40[a]
	$\Delta\delta$	-0.32 ± 0.32	-0.34 ± 0.53	-0.20 ± 0.44	0.11 ± 0.30[a]
GOODS	$\Delta\alpha$	-0.20 ± 0.23	0.15 ± 0.42[a]	0.03 ± 0.44[a]	-0.10 ± 0.39[a]
	$\Delta\delta$	0.22 ± 0.28	-0.18 ± 0.37[a]	0.03 ± 0.44[a]	-0.18 ± 0.33[a]

Column headers− (1) Miller et al. (2013), (2) Kellerman et al. (2008) (3) Norris et al. (2006), (4) Afonso et al. (2006)

Notes− (a) Few ($\leqslant 15$) sources available to calculate the offsets.
(b) Positional uncertainty not mentioned in catalog; assumed to be $0.1''$.
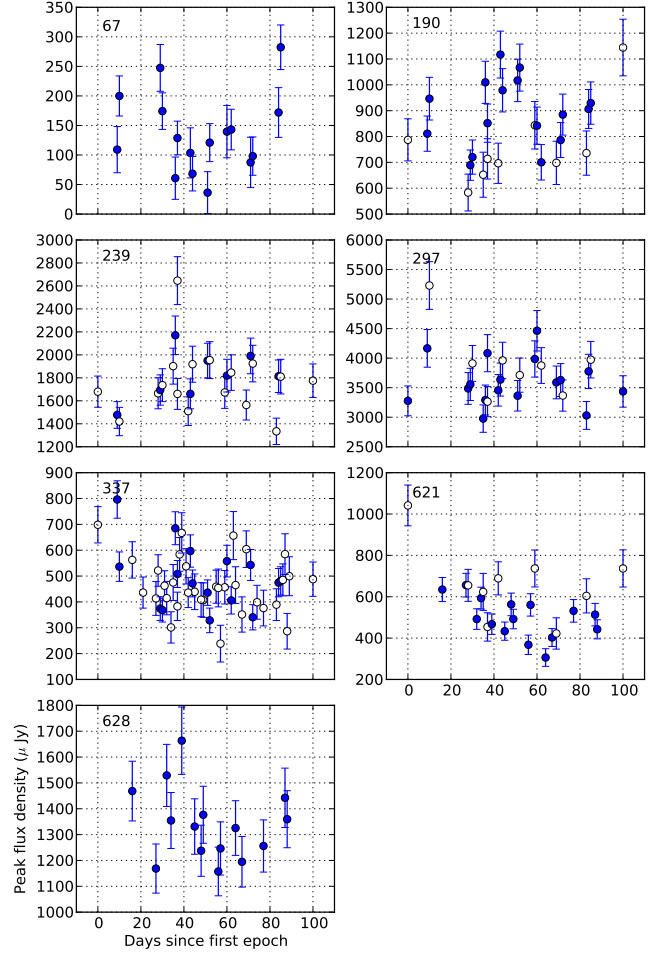(c) All offsets are in arcseconds.

FIG. 5.— Light curves of the variable sources from Miller et al. DR2 catalog listed in Table 3. Filled symbols represent the flux densities considered for variability analysis (i.e. where $\theta < 12'$). Flux densities in epochs where a source lies within the 50% power circle of the beam (i.e. where $\theta \lesssim 15'$) are plotted for reference as open symbols. The error bars take into account the background rms, primary beam correction, and bandwidth smearing (no taper).

3.3. *Notes on Variables Found*

An important question that we can address is whether the variability at sub-milliJansky levels is dominated by normal star-forming galaxies or by AGN. For AGN-dominated samples above 1 mJy, variability at frequencies of a few GHz or below is thought to be dominated by propagation effects (i.e. refractive interstellar scintillation) and not by intrinsic changes in the source (Gaensler & Hunstead 2000; Ofek & Frail 2011). For a disk galaxy we expect there to be steady emission from diffuse synchrotron emission and the sum of all supernovae. Intrinsic variability can be induced by (i) a stellar explosion (supernova, low-luminosity gamma-ray burst), and (ii) nuclear radio emission (AGN). Mapping the radio emission to the center of the galaxy would favor (ii) whereas if the radio emission is mapped to the disk then (i) is favored.

To this end, we overplotted radio source positions of the seven variables given in Table 3 on HST image cutouts from GOODS-S and the GEMS projects (Figure 6). All of the radio sources have an optical counterpart on these HST images. The redshifts of these
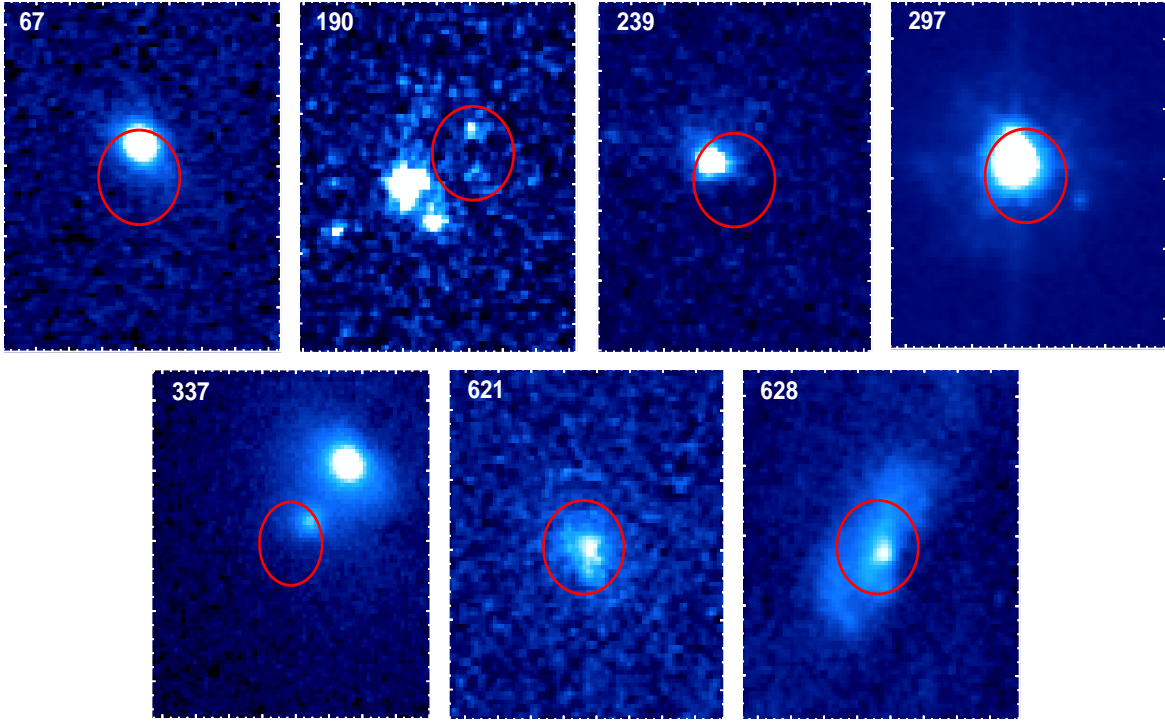
FIG. 6.— $2.2'' \times 2.5''$ GEMS (all sources except ID 337) and GOODS-S (ID 337) F606W ACS-WFC image cutouts of variable sources found in this work. The red error ellipses denote the radio source positions from Miller et al. (2013), shift-corrected to the HST source positions. All the positional uncertainties are 1σ (see § 3.2 and Table 4).

objects vary from 0.3 to 2.3. Thus, the post-offset radio source-position uncertainties lie between 1.4 and 2.7 kpc. Within most error ellipses, there is a centrally compact source, suggesting that this is the source of the variable emission. Light-curves of the variable sources are given in Figure 5. Detailed notes on each of these objects are given below, and key physical parameters are listed in Table 3. All of the variable sources have luminosities in excess of 10^{30} erg s^{-1} Hz^{-1}, where the luminosity functions of AGNs and star-forming galaxies intersect (Condon 2002). Most of the variable sources have a roughly flat spectrum between 1.4 GHz and 5.5 GHz suggesting the presence of AGN. Note that the spectral indices between 1.4 GHz and 5.5 GHz are based on non-simultaneous measurements having different resolutions. Additionally, mid-infrared colors, far-infrared spectral indices, and mid-to-far-infrared luminosities indicate that most of the variable sources reside in star-forming galaxies. In Figure 7 we show the mid-infrared color-color diagram for the variable sources using *Spitzer*/IRAC photometry from the SIMPLE survey catalog (Damen et al. 2011). The Donley et al. (2012) AGN selection region is overplotted. For comparison, the colors of the radio sources from the AEGIS20 sample (Willner et al. 2012) are shown along with those of three template spectral energy distributions (SEDs) — an elliptical galaxy, an Sbc galaxy, and an AGN — from Assef et al. (2010). This figure along with the far-infrared flux densities identify two variables with AGNs and the rest as star-forming galaxies.

Summarizing, the high-resolution optical images together with photometric information from radio and mid-to-far-infrared suggest that variability arises from the central regions of an AGN or star-forming galaxy.
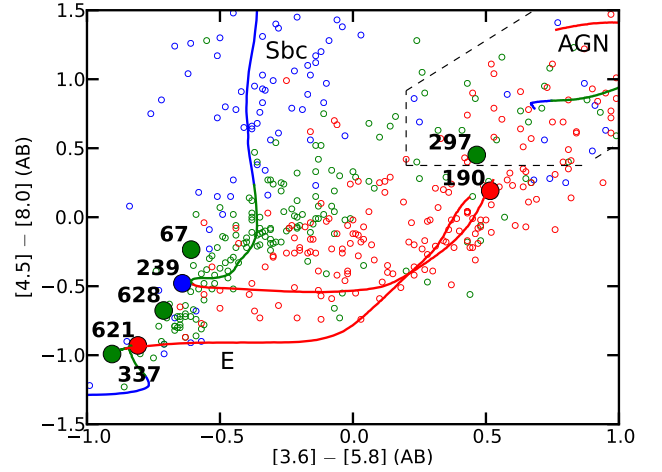

FIG. 7.— *Spitzer*/IRAC color-color diagram for the mid-infrared counterparts of the variable sources (filled circles). For comparison, the radio sources from AEGIS20 (Willner et al. 2012) are also shown (unfilled circles). The AGN selection region in the upper-right corner, bounded by the dashed lines, is from Donley et al. (2012). Curves represent the colors of three template spectral energy distributions (E: elliptical galaxy; Sbc: spiral galaxy; AGN) from Assef et al. (2010) as redshift increases from 0 to 3. The redshift-dependent color-coding is — blue: $z \leqslant 0.5$, green: $0.5 < z \leqslant 1.1$ and red: $z > 1.1$.

ID 67: This source has a mean 1.4 GHz flux density of 136 ± 9 μJy and shows the strongest variation among all seven variables. Two flaring bursts are evident from the light curve (Figure 5), one of them lasting for about 25 days, and the flux density of the other increasing by

at least a factor of 2.5 in 12 days. This source also appears in the Kellerman et al. 1.4 GHz catalog, where its flux density is 90 ± 16 μJy. The photometric redshift of z=1.005 (Bonzini et al. 2012) implies a mean radio luminosity of log $L_R = 30.6$ erg s^{-1} Hz^{-1}. The radio spectral index (defined as α where $S_\nu \propto \nu^\alpha$) between 1.4 GHz and 5.5 GHz is +0.23 (Huynh et al. 2012). The morphology of the host galaxy as seen from the HST image appears to be that of a bright, compact nucleus surrounded by faint extended structure. The half-light radius of the galaxy according to the *GALFIT* parameters from the GEMS catalog is 620 ± 10 pc. The K-corrected[10] mid-to-far-infrared spectral luminosities derived from the FI-DEL and GOODS (Magnelli et al. 2009, 2011), and SIMPLE *Spitzer* surveys are $\nu L_\nu (5.8$ μm$) = 6.6 \times 10^8$ L$_\odot$ and $\nu L_\nu (24$ μm$) < 5.4 \times 10^{10}$ L$_\odot$. By comparing these quantities with the Chary & Elbaz (2001) template SEDs (see Figure 4 of that paper) and from the *Spitzer*/IRAC color-color diagram (Figure 7), we interpret that the host galaxy is star-forming.

Taken together, the radio luminosity, radio spectral index, optical morphology and mid-to-far-infrared flux densities argue that ID 67 is a star-forming galaxy harboring a low-luminosity AGN.

ID 190: This source has a mean 1.4 GHz flux density of 891 ± 12 μJy. The light curve appears to fluctuate between high and low flux density states on a timescale of order 20 days. This source is also found in the Kellerman et al. and Norris et al. catalogs where the flux density is 970 ± 30 and 810 ± 19 μJy respectively, and a spectral index of -0.6 between 1.4 and 4.8 GHz. The 5.5 GHz flux density from Huynh et al. is 555 ± 17 μJy, for which we derive a spectral index of -0.35. The photometric redshift of z=2.296 (Bonzini et al. 2012) implies a mean radio luminosity of log $L_R = 32.1$ (cgs). The HST image of this galaxy as seen in Figure 6 shows that the radio source is offset from the brightest emission in the field. On the basis of its departure from the radio-FIR correlation, Norris et al. classify this source as an AGN. The radio and far-infrared spectral indices (Table 3) and mid-infrared colors (Figure 7) are consistent with this identification. Note that between 24 μm and 70 μm, a spectral index greater than -1.5 is representative of AGNs (e.g. Condon 2002).

ID 239: This source has a mean 1.4 GHz flux density of 1822 ± 27 μJy. This source is also found in the Kellerman et al. and Norris et al. catalogs where the flux density is 2030 ± 43 and 1640 ± 20 μJy respectively. Its light-curve shows a gradual increase in flux density, over a period of about 50 days, followed by an equally gradual decline. The radio spectral index between 1.4 GHz and 5.5 GHz is -0.30. This slope is more consistent with AGNs than star-forming galaxies which are expected to have $\alpha \simeq -0.8$. The redshift z=0.266 (Treister et al. 2009) implies a mean radio luminosity of log $L_R = 30.5$ (cgs). The HST cutout (Figure 6) reveals just a bright nucleus surrounded by faint extended structure, the half-light radius being 3.2 ± 0.5 kpc. The K-corrected mid-to-far-infrared spectral luminosities, $\nu L_\nu (5.8$ μm$) = 3.6 \times 10^8$ L$_\odot$ and $\nu L_\nu (24$ μm$) < 1.0 \times 10^9$ L$_\odot$, and the mid-infrared colors

(Figure 7) show that the host is a star-forming galaxy.

Taken together, the radio luminosity, radio spectral index, optical morphology and mid-to-far-infrared flux densities argue that ID 239 is a star-forming galaxy harboring a low-luminosity AGN.

ID 297: This source has a mean 1.4 GHz flux density of 3.60 ± 0.04 mJy. Huynh et al. (2012) measure a 5.5 GHz flux density of 12.25 mJy, implying a spectral index of +0.89. One the basis of this steep positive spectral index, Huynh et al. (2012) suggest that this is part of a class of Gigahertz Peaked Spectrum (GPS) sources, thought to be a young AGN. The redshift z=0.605 (Treister et al. 2009) implies a mean radio luminosity of log $L_R = 31.5$ (cgs). ID 297 appears to be a stochastically varying source with no specific trend in its light-curve (Figure 5). The radio position is consistent with a bright, unresolved HST source. On the basis of departure from radio-FIR correlation Norris et al. classify this source as an AGN. The mid-infrared colors (Figure 7) and far-infrared spectral index (Table 3) are consistent with this identification.

ID 337: This source has a mean 1.4 GHz flux density of 494 ± 11 μJy. From its radio light-curve, ID 337 appears to have repeated outbursts roughly every 25d. The flux density of this source in the Miller et al., Kellerman et al., Norris et al. and Afonso et al. catalogs is 439 ± 8, 524 ± 14, 380 ± 16, and 404 ± 34 μJy respectively, and a spectral index of -0.2 between 1.4 and 4.8 GHz. The 5.5 GHz flux density from Huynh et al. is 443 ± 20 μJy, for which we derive a spectral index of -0.08. The photometric redshift of z=0.734 (Vanzella et al. 2008) implies a mean radio luminosity of log $L_R = 30.8$ (cgs).

The HST image (Figure 6) shows that the optical counterpart to the radio source is actually associated with the fainter (or more extincted) galaxy among a group of two closely separated galaxies. This has lead to some confusion over the correct optical identification. VLA observations carried out in 1999–2001 and 2007 (Kellerman et al. and Miller et al. respectively) suggest that the radio counterpart is the fainter galaxy. ATCA observations carried out in 2003–2004 (Norris et al. and Afonso et al.) suggest the bright galaxy as the optical counterpart. Our optical-radio frame tie summarized in Table 4 supports the fainter optical source as the likely radio counterpart. The spectral index and radio luminosity argue that ID 337 is a low-luminosity AGN. Afonso et al. state that the bright radio source is a luminous star-forming galaxy, possibly part of a merging system. Comparison of the K-corrected spectral luminosities, $\nu L_\nu (5.8$ μm$) = 4.0 \times 10^9$ L$_\odot$ and $\nu L_\nu (24$ μm$) < 1.1 \times 10^{10}$ L$_\odot$, with the Chary & Elbaz (2001) SED templates and Desai et al. (2007), together with the mid-infrared colors (Figure 7) advocate the star-forming nature of the host galaxy.

ID 621: This source has a mean 1.4 GHz flux density of 497 ± 10 μJy. The first epoch in its light-curve reveals an initial brightening of the source followed by a decline and subsequent small-amplitude variations. The maximum flux density is close to 400% of the quiescent flux density of about 300 μJy. The Miller et al., Kellerman et al., and Norris et al. catalogs list the flux density of this source as 494 ± 10, 565 ± 17, 450 ± 18 μJy respectively, and a spectral index of -0.1 between 1.4 and 4.8 GHz. The

---

[10] For the 5.8 μm and 24 μm luminosities, the K-correction has been applied based on the 8 μm and 70 μm flux densities. In the absence of a 70 μm detection, the 3σ upper limit has been used.

5.5 GHz flux density from Huynh et al. is 689 ± 16 μJy, for which we derive a spectral index of 0.24. The photometric redshift of z=1.107 (Bonzini et al. 2012) implies a mean radio luminosity of log $L_R = 31.2$ (cgs).

The HST image reveals a bright nucleus surrounded by diffuse emission. The half-light radius is $\sim$3 kpc. The radio position is consistent with the nuclear source. Taken together, the radio luminosity, spectral index and optical morphology argue that ID 621 is an AGN. Additionally, the K-corrected 5.8 μm luminosity of 5.0×10^9 L$_\odot$, the 24 μm luminosity upper limit of 6.8×10^{10} L$_\odot$, and the mid-infrared colors (Figure 7) suggest that this is also a star-forming galaxy.

ID 628: This source has a mean 1.4 GHz flux density of 1.34 ± 0.02 mJy. The light-curve indicates a steady increase of the quiescent emission to maximum flux density, followed by a steady decline. The Miller et al., Kellerman et al., and Norris et al. catalogs list the flux density of this source as 1.07 ± 0.02, 1.33 ± 0.03, and 0.90 ± 0.02 μJy. The 5.5 GHz flux density from Huynh et al. is 0.78 ± 0.02 μJy, for which we derive a spectral index of -0.40. The photometric redshift of z=0.685 (Norris et al. 2006) implies a mean radio luminosity of log $L_R = 31.2$ (cgs).

The HST image reveals a bright nucleus surrounded by disk-like diffuse emission, the half-light radius being $\sim$10 kpc. The radio position is consistent with the nuclear source. On the basis of departure from radio-FIR correlation Norris et al. classify this source as an AGN, which is consistent with our flat radio spectral index and radio luminosity. Additionally, the K-corrected 5.8 μm luminosity of 3.0×10^9 L$_\odot$, the 24 μm luminosity of 3.7×10^{10} L$_\odot$, and the mid-infrared colors (Figure 7) suggest that this is a normal star-forming or a starburst galaxy. Taken together, the radio and mid-to-far-infrared properties along with the optical morphology argue that ID 628 is an AGN embedded within a star-forming galaxy.

4. TRANSIENT SEARCH

For our transient search, we are interested in identifying those point sources which show up above the flux density limit for a short amount of time (corresponding to one or more epochs depending on the cadence), and remain below the limit in all other epochs. It is therefore important to reliably distinguish noise from real transients (which will determine how well we can reject false positives), and to avoid the rejection of transients as noise (reducing the number of true negatives). Thus, characterizing the effectiveness of source-finding algorithms in terms of reliability and completeness is crucial. Another motivation for characterizing source-finding algorithms is to find the optimum parameter values for use in transient searches.

In the following subsections we begin by the testing the efficacy of existing source-finding algorithms (§4.1). We then apply some of the better-performing algorithms on the E-CDFS epochs (§4.2).

4.1. *Efficacy of Source-finding Algorithms*

Recently, quantitative tests have been carried out on the reliability and completeness of source-finding algorithms (Huynh et al. 2011; Hancock et al. 2012). Of these many publicly available software packages, `sfind`

and `IMSAD` in *MIRIAD*, *Aegean*, *SExtractor*, and *Selavy*, these studies found that `sfind` and *Aegean* produce the most reliable catalogs.

The analysis of Huynh et al. (2011) and Hancock et al. (2012) was carried out on two simulated data sets. (i) The ASKAP simulation is a $4° \times 4°$ image of a full continuum observation with critically-sampled beams in the 6-km ASKAP configuration. Its pixel-scale is $2.75''$ and the rms noise is about 35 mJy, which varies across the field. 16 idealized beams one degree apart, spaced in a rectangular grid, mimic the effect of the phased-array feed. The image contains $\sim$7.7 million sources having flux densities greater than 1 μJy from the S^3-semi-empirical extragalactic simulation (Wilman et al. 2010). (ii) For the Hancock et al. simulation, a sky image was created as a 4801^2 pixel image $8°$ across with a $6''$ pixels sampling a $30''$ beam and an rms noise of 25 μJy. Sources were injected at random positions with angular sizes (with random position angles) from $0''$ to $52''$ and with source number counts distributed with peak flux densities as N(S)$\propto$ S$^{-2.3}$ such that 15,000 sources having fluxes densities $> 1\sigma$ are present in the image.

With our E-CDFS dataset we are able to carry out a similar analysis using real data with all its attendant residual calibration and imaging errors. A comparison of real and simulated data could be informative. While simulations are useful in determining which source-finding algorithm work best in general, they do not explore the parameter space of the algorithm thoroughly. Hence they may not provide optimum parameter values for a transient search on a specific dataset. The deep field of the E-CDFS is well suited for this comparison. The 4096^2 pixel image is $34'$ across with a $0.5''$ pixels sampling a synthesized beam of $2.8'' \times 1.6''$ (position angle$\simeq 0°$) and an rms noise 7.4 μJy. As outlined in §3, great care was taken in constructing the DR2 source catalog so we can be assured of its completeness and reliability (see also Figure 2).

In what follows, we will use the DR2 catalog and the deep E-CDFS image to test various source-finding algorithms for completeness and reliability for different input parameters. We use the terms "real sources" and "false sources" as being those sources present in the DR2 catalog, and those that are not, respectively[11]. Whether a source detected by a source-finding algorithm has a counterpart in the DR2 catalog is determined by searching for DR2 sources within $1''$ from the source position. Following Huynh et al. (2011), we define two additional terms, "completeness", as the fraction of real sources detected by a source-finding algorithm, and "reliability", as the fraction of detected sources which are real. Note that in Figures 8, 9 and 12, we plot the completeness and reliability within contiguous SNR bins, unlike Huynh et al. (2011) and Hancock et al. (2012) where, for a given SNR, the plots represent the completeness and reliability for sources greater than or equal to that SNR. For source-finding algorithms employing a probabilistic approach of drawing a pixel from the background and thus

---

[11] Even though the DR2 catalog was constructed with great care, it is likely that $\sim$1 genuine source was missed and a handful of spurious sources added (perhaps not truly spurious, but `SAD` sources at 4σ bumped up to 5σ by the `JMFIT` task in *AIPS*). Recall that according to Eddington bias, more sub-5σ sources get bumped up than 5σ get bumped down.

calculating false-detection rate (FDR), FDR + Reliability = 100%.

4.1.1. `sfind` (MIRIAD)

In its default mode, `sfind`[12] incorporates a statistically robust method for detecting source pixels, called "False Discovery Rate", or FDR. In the FDR algorithm, detected sources are drawn from a distribution of pixels with a robustly known chance of being falsely drawn from the background. Contiguous, monotonically decreasing adjacent pixels from the FDR-selected ones, are used for fitting 2-D elliptical Gaussians to the sources. Thus, the fraction of expected false sources is more reliably determined than in sigma-clipping criteria (see methods below). Details of the FDR method can be found in Hopkins et al. (2001). The run-time for `sfind` searching for sources down to 5σ in the DR2 image is about 20 seconds. For all the tests carried out on `sfind`, MIRIAD version 4.2.3 (optimized for CARMA; CVS Revision 1.11, 2011/04/26) was used.

Here, we explore the completeness and reliability of `sfind` by tweaking the two relevant parameters: (i) *alpha* (α), the percentage of probable background pixels that can be accepted in the analysis, and (ii) *rmsbox*, the size of the smoothing box used for estimating the background and the standard deviation of the image. Table 5 lists the parameter values tested.

In general, completeness of `sfind` increases, and its reliability decreases (Table 5) with increasing α, as expected for an FDR algorithm. Also, both these quantities increase with *rmsbox*. The rise in completeness and reliability is precipitous (between 5% and 20%) as *rmsbox* is increased from 5 to 10 beamwidths, after which it flattens off. There is a slight decrease in the reliability as *rmsbox* is increased from 20 to 50 beamwidths. 10 beamwidths can then be interpreted as the minimum box size for determining the background rms noise without significant contribution from the sources themselves, whereas 50 beamwidths would correspond to the size where the calculated rms starts deviating significantly from the true local rms. However, in the case of bright sources in the field, these *rmsbox* limits might be somewhat larger. Completeness and reliability are $\sim$88% for α=10 and *rmsbox*=20, but better reliability (by a few percent) can be obtained at an equal expense of completeness by using α=5 or 2. Further decrease in α substantially decreases the completeness without any significant improvement in the reliability. Thus, we determine the optimum values for the input parameters to be $5 \leqslant \alpha \leqslant 10$ and $10 \leqslant rmsbox \leqslant 50$ beamwidths. For the FDR algorithm, we expect the reliability to be $1 - \alpha$. However, `sfind` reliability is lesser than this expected value by a few percent, due to the acceptance of sidelobes of bright sources and fitting of extended sources with multiple elongated and overlapping (unphysical) components. To some extent, the acceptance of these false sources can be reduced by choosing a lower value of α. Some examples of the components fit by `sfind` to different kinds of sources in different environments in the DR2 image are shown in Figure 6. In general, compared to other source-finding algorithms excepting `SAD`, `sfind` has a better reliability for a given completeness, and vice-versa (Tables 5–9; see

also §4.1.3 for a note on the efficacy of `SAD`). These findings are consistent with the tests carried out on simulated images (Huynh et al. 2011). In the ASKAP simulation, the reliability of the `sfind` catalog is rather flat with respect to α, whereas, for the Hancock et al. simulation it decreases by several percent as α increases from 0.1 to 10 (Huynh et al. 2011). In the latter, *rmsbox*=20 is found to give slightly better reliability than *rmsbox*=20. For the $\sim$3–10σ sources in these simulations, the completeness for α=5 is greater than that for α=0.1 by 5–10%. They do not explore the completeness for α=10 and *rmsbox* other than 10 beamwidths.

Figure 8 shows the completeness of the `sfind` catalog as a function of the SNR of the detected sources for different values of the input parameter α. The reduction in completeness beyond SNR$\sim$20 is due to missed or badly-fit components of extended sources. The completeness for optimum values of input parameters is shown in the upper panel of Figure 9, which shows that `sfind` reaches 100% completeness at a much lower SNR than other algorithms. The lower panel of Figure 9 shows the reliability. Curiously, there is a dip in the reliability between SNRs of 6 and 10. Point sources strewn across the DR2 image but primarily located near its edges (in regions of increased rms; see the lower panel of Figure 10 for example), which are absent in the DR2 catalog and detected by `sfind`, are responsible for this reduced reliability. Some of these sources get rejected from the `sfind` catalog when a lower value of α or *rmsbox* is used, indicating that they are either not genuine sources on the sky or are sources at a lower SNR. Indeed, other algorithms find some of these sources to be at a much reduced SNR between 3 and 5. The results of (Hancock et al. 2012) also show a dip in the `sfind` reliability, although centered on SNR$\sim$15. Figure 11 compares the SNR of sources detected by `sfind`, with those of the counterparts from the DR2 catalog. The largely increased SNR reported for sources at or below 6σ in the DR2 catalog, and largely reduced SNR for sources above $\sim$100σ, is anomalous. Although for SNR above 100 the discrepancy is likely to be due to overlapping components fit to extended sources, all the above observations indicate that the SNR reported by `sfind` is somewhat different from that reported by other algorithms. Sources with 6<SNR<10 in the `sfind` catalog clearly have over-estimatated (upto 100%) peak flux densities and different dimensions than what would be expected by inspecting the cutouts of these sources. Thus, fitting of sources rather than differently-reported rms is the cause of the SNR descrepancy, at least at these low SNRs. The solution of this issue is possibly in correctly setting the *fdrpeak* and *psfsize* input parameters which allow reasonable measurements of sources close to the threshold. Since there are not many sources in our data having SNR $\gtrsim$70 (Figure 2), the suggested trend in completeness, reliability and measured SNR in this domain should be treated with caution.

4.1.2. *SExtractor*

SExtractor[13] is a source-finding program widely used in optical astronomy, and is particularly oriented towards the reduction of large-scale galaxy-survey data, as well

---

[12] www.atnf.csiro.au/computing/software/miriad/doc/sfind.html

[13] www.astromatic.net/software/sextractor

as sparsely- to moderately-crowded stellar fields. It analyzes the image in two passes such that, in the first pass, a background map is made, and in the second, background-subtraction, filtering and thresholding is done on-the-fly. Detected sources are then deblended and CLEANed before performing photometry. The run-time for *SExtractor* searching for sources down to 5σ in the DR2 image is about 2 seconds. For all the tests carried out here, the latest release of *SExtractor* (version 2.8.6, 09-Apr-2009) has been used.

We adopted a strategy of searching for sources down to $N\sigma$ (N=3,5,6,7,10, as set using the $DETECT\_THRESH$ parameter in *SExtractor*) and then selecting the ones $\geqslant 5\sigma$ prior to the comparison with the DR2 sources. Two input parameters (apart from possibly $DEBLEND\_NTHRESH$, $BACK\_FILTTHRESH$ and $CLEAN\_PARAM$, which we have not tested) are expected to affect the completeness and reliability of the *SExtractor* catalog. First, the mesh-size, which determines the size of the box used for background rms estimation, is specified by $BACK\_SIZE$, and second, the size of the median filter ($BACK\_FILTERSIZE$) applied to the background grid used for smoothing large artifacts in the image.

In Table 6, we list the completeness and reliability as a function of these parameters. Note that in this table, in order to reflect the true completeness for the specified detection threshold, the completeness values of the *SExtractor* catalog for 6σ, 7σ and 10σ thresholds have been normalized by the percentage of sources in the DR2 catalog which are beyond these thresholds respectively. Thus, the completeness at any threshold is expected to be 100% if all the DR2 sources beyond that threshold are detected by *SExtractor*. The general trend observed is that, with increasing detection threshold, the completeness of *SExtractor* is fairly steady, but its reliability increases. However, for 5σ and 6σ detection thresholds, the completeness is rather low ($\sim$75% or lower). Choosing higher detection thresholds such as 7–10σ does not seem to improve the completeness substantially. This is due several factors through which the source-finding in *SExtractor* seems to be different from the conventionally-used algorithms in radio astronomy. Firstly, for any given source, the reported peak flux density decreases as the detection threshold is raised, although there is not much of a change in the reported rms. Hence, the SNR of sources close to the detection threshold decreases so as to be rejected by *SExtractor*. Choosing a 3σ threshold usually gives the correct peak flux density. This is also the reason why searching down to 3σ and selecting sources at the desired, higher threshold, increases the completeness (although with the side effect of reduced reliability). The explanation for such an effect is hinted by the fact that the dimensions of the fitted sources decrease (this effect is quite significant for sources with SNR<10) with increasing detection threshold, implying that the number of source pixels considered in the fitting process depends on the threshold. Secondly, several DR2 sources having SNR<10 and detected by other source-finding algorithms (excepting IMSAD), are not detected at all by *SExtractor*, not even at a reduced SNR. This might be due the differences in the fitting process rather than discrepant rms. Lastly, for some extended sources, positions of the components reported by *SExtractor* are at

least a few synthesized beams away from the positions given in the DR2 catalog; the latter almost always agree with the positions reported by other sources-finding algorithms (again, with the exception of IMSAD; see §4.1.4 for anomalies associated with IMSAD). The completeness peaks at a background mesh-size of 5 beamwidths used with 3x smoothing, whereas reliability appears to increase steadily or remain constant with increasing mesh-size. $BACK\_FILTERSIZE$ is not seen to change either the completeness or reliability much. *SExtractor* does not seem to achieve a high completeness and reliability (>85%) simultaneously for any given set of values for the input parameters. For reliability of > 90%, we see that completeness < 80%. Beyond a detection threshold of 6σ, the reliability is quite good, yet less than 100% owing primarily to the acceptance of sidelobes and the fitting of extended sources handled differently. Nevertheless, the highest reliability for *SExtractor* is better than the catalogs of other algorithms (except, maybe SAD), and hence it is best for cases where reliability is strongly favored over completeness. In such a case, we recommend setting the detection threshold to about 10σ, along with a large enough value for $BACK\_SIZE$. Tests on *SExtractor* with the ASKAP and Hancock et al. simulations (Huynh et al. 2011) suggest that the reliability is almost constant for mesh-size between 10 and 100 beamwidths, but increases with the detection threshold. They find that the completeness generally decreases or remains constant as mesh-size increases, but its change with respect to detection threshold is not explored. Our results are thus broadly consistent with the tests on the simulated images.

The upper-left panel of Figure 8 shows the completeness of *SExtractor* sources as a function of their SNR in the DR2 catalog for mesh-size of 20 beamwidths and different values of the detection threshold. The completeness for optimum input parameter values is shown in the upper panel of Figure 9. *SExtractor* approaches 100% completeness faster than IMSAD, but slower than the other algorithms. This is possibly related to the approach used for determining source pixels. The decrease in completeness beyond SNR$\sim$20 can be attributed to the extended-source components being reported differently than other algorithms, as mentioned above. The lower panel of Figure 9, which shows the *SExtractor* reliability in relation with other source-finding algorithms, reveals that this algorithm gives quite unreliable results for SNR$\lesssim$10. These low-SNR false sources are all point sources located near the edges of the DR2 image, where the rms is somewhat large (12 μJy), but the rms reported by *SExtractor* is quite small (7–8 μJy). This scenario is presented in the image cutouts in Figure 10, which also illustrate the ability of *SExtractor* to find sources with different morphologies located in different environments. Figure 11 shows the SNR of sources detected by *SExtractor*, in comparison with the corresponding sources in the DR2 catalog. For SNR$\leqslant$30, the peak flux density of *SExtractor* sources sfind,SExtractor is generally less than that of the DR2 catalog sources, suggesting that the difference in the calculated peak flux density (and to a smaller extent, the associated uncertainty) is responsible for the observed departure of the *SExtractor* SNR from the DR2 SNR.
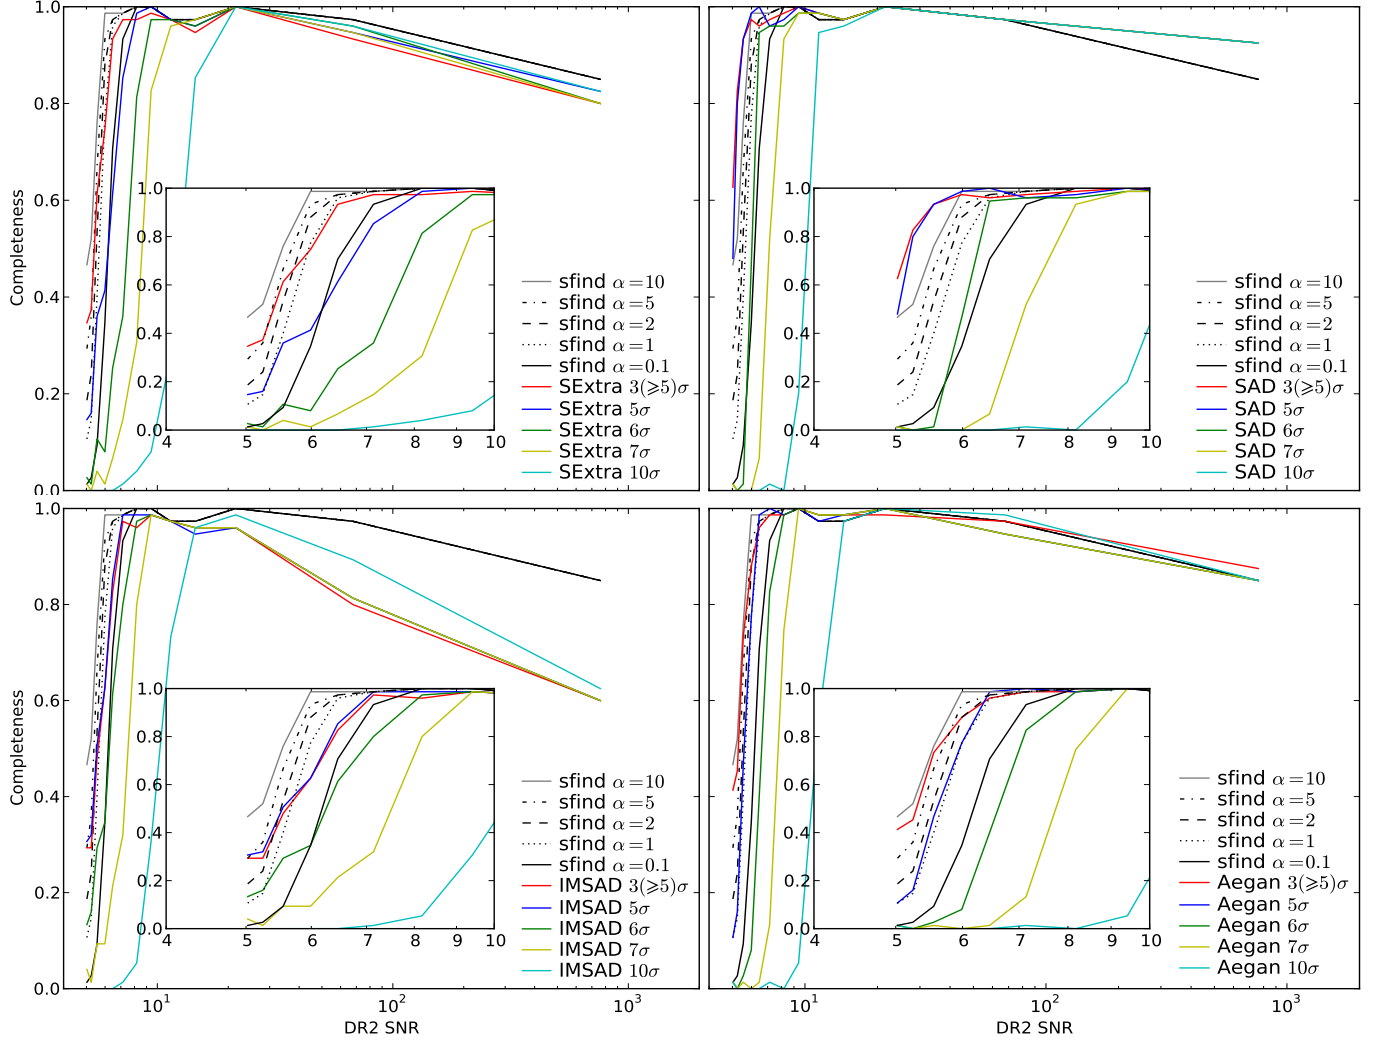
FIG. 8.— Completeness of the catalogs generated by different source-finding algorithms. Results for a background mesh-size (rmsbox) of 20 beamwidths, wherever specifiable, are shown. For IMSAD and *Aegean*, the results are for the histogram option and csigma=$1\sigma_{\mathrm{cmap}}$ respectively. See §4.1 for the definition of completeness used here. The decreasing completeness beyond $SNR \sim 50$ is due to deblending of multiple components of extended sources. The inset shows only the region where SNR is between 4 and 10. Smoothing over every 75 data points has been done before plotting. Note the low-number statistics for sources with SNR $\gtrsim 70$ as implied by Figure 2.

4.1.3. SAD *(AIPS)*

The "Search and Destroy" (SAD[14]) algorithm finds all the pixels above a specified threshold (typically a multiple of the rms noise, which is assumed to be Gaussian) in the image, and merge contiguous pixels above the threshold into islands. The strength and size of each island is then estimated, followed by least squares Gaussian fitting of each island (if rms residual is too high, then multiple Gaussian fits may be applied). However, note that Gaussian statistics may not be a good model for the distribution of values in pixels well above zero flux density on account of thermal noise, and calibration and imaging artifacts (Cotton & Peters 2011). The runtime for SAD searching for sources down to 5σ in the DR2 image is about 2 minutes. Additionally, to prepare the background rms image, the task RMSD takes ~ 15 minutes. For all the tests carried out on SAD, 31DEC11 *AIPS* was used.

[14] www.aips.nrao.edu/cgi-bin/ZXHLP2.PL?SAD

SAD has several input parameters which affect the number of sources detected. Here, we test the effect of the detection SNR ($CPARM$), rms threshold, flux residual threshold ($DPARM, GAIN$), and the size of rms box ($IMSIZE$). The parameter $DPARM(3)$ (along with $GAIN$ added in quadrature) specifies the upper limit for the rms in the fitting box, while $DPARM(7)$ and $GAIN$ specify the upper limit on the residual flux in the fitting box. $GAIN$ thus defines the fraction of the source flux which is acceptable in the residual image. By default, SAD uses the entire image to find the rms. We used the *AIPS* task RMSD to prepare rms images using mesh-sizes (specified by the $IMSIZE$ parameter) of 5, 10, 20, and 50 beamwidths. Although a decremental search in SNR (via $CPARM$) is recommended in the SAD help file, we found that such a search results in multiple sources being fit to a single genuine source during each iteration, especially when the source is extended. Hence, we rejected this recommendation. As in the case of *SExtractor*, we searched for sources down to Nσ. Note that in SAD, the errors in the flux density are determined by theory from

the image rms (`actnoise` keyword in the image header).

In Table 7, we list the completeness and reliability of the `SAD` catalog for different values of the input parameters. We held the input parameters $DPARM(3)$ and $DPARM(7)$ fixed at 1000 and 1 (in units of Jy beam^{-1}) respectively in order to get optimum completeness (without significant loss of reliability, as we found later, but possibly at the expense of correct SNR of the detected sources). Due to these large input values, we did not find any change in the results with the $GAIN$ parameter. However, we found that reliability can be traded for completeness by setting smaller values for $DPARM(7)$ and $GAIN$. In general, the completeness and reliability of `SAD` increase with the detection threshold. Completeness increases with mesh-size. So does reliability, though this quantity decreases significantly as the mesh-size is increased from 20 to 50 beamwidths, similar to `sfind`. Our inspection of the rms image for a mesh-size of 50 beamwidths reveals that there are a few pockets where the rms is rather low (~ 1 μJy). A profusion of false sources (or very low-SNR sources reported to have an SNR above the threshold) detected in these pockets are responsible for the markedly reduced reliability for the case of mesh-size equal to 50 beamwidths. Search with `actnoise` usually performs at least a few percent worse in terms of completeness and reliability than using mesh-size of 20 beamwidths. Curiously, the completeness and reliability of `SAD` is >90% for a wide range of input parameters tested. We determine the optimum values of input parameters to be 10–20 beamwidths mesh-size and 6–7σ detection threshold. Searching down to 3σ followed by selection of sources greater than 6–7σ may improve completeness to some extent. Huynh et al. (2011) and Hancock et al. (2012) have not tested `SAD` on simulated images.

The completeness of `SAD` with respect to detection SNR and for different detection thresholds is shown in the upper-right panel of Figure 8, which depicts the high level of completeness close to the threshold and rapid increase with SNR compared to the other algorithms. The upper panel of Figure 9 plots the `SAD` completeness for optimum input parameters. Due to missed sources throughout the SNR spanned, the completeness is seen to hover close to unity, but not quite getting to 100%. The source-rejection criteria based on 2-D Gaussian fitting in `SAD`, defined by the several elements of the $DPARM$ input array, is the likely cause for such such missed sources. The lower panel of Figure 9 shows the reliability for the optimum values input parameters. As with the case of completeness, the reliability seems to hover close to unity. `SAD` provides the best reliability for sources SNR$\lesssim 20$, beyond which, `sfind` gives better results. Departure of the reliability from unity for a wide range of SNR values is due to the fitting of different components to extended sources. Some examples of how `SAD` fits different kinds of sources in the DR2 image, in comparison with other algorithms, is shown in Figure 10. Figure 11 shows the SNR of sources detected by `SAD`, in comparison with that of their counterparts in the DR2 catalog. Usually the SNRs agree with each other, but any disagreement between the two is due to the reported uncertainty in the peak. The peak flux densities reported by `SAD` match those in the DR2 catalog quite well.

There is a possibility that the superior completeness and reliability of `SAD` may be due simply because the DR2 catalog, against which we are doing all of our comparisons, was constructed from `SAD`. However, the construction of the DR2 catalog involved much more than running `SAD` (see §3 for more details). In any case, an independent check using simulated sources is warranted. In Figure 12, we compare the completeness and reliability of `SAD` with other source-finding algorithms on the Hancock et al. (2012) simulated image. We find a similar superior behavior of `SAD`.
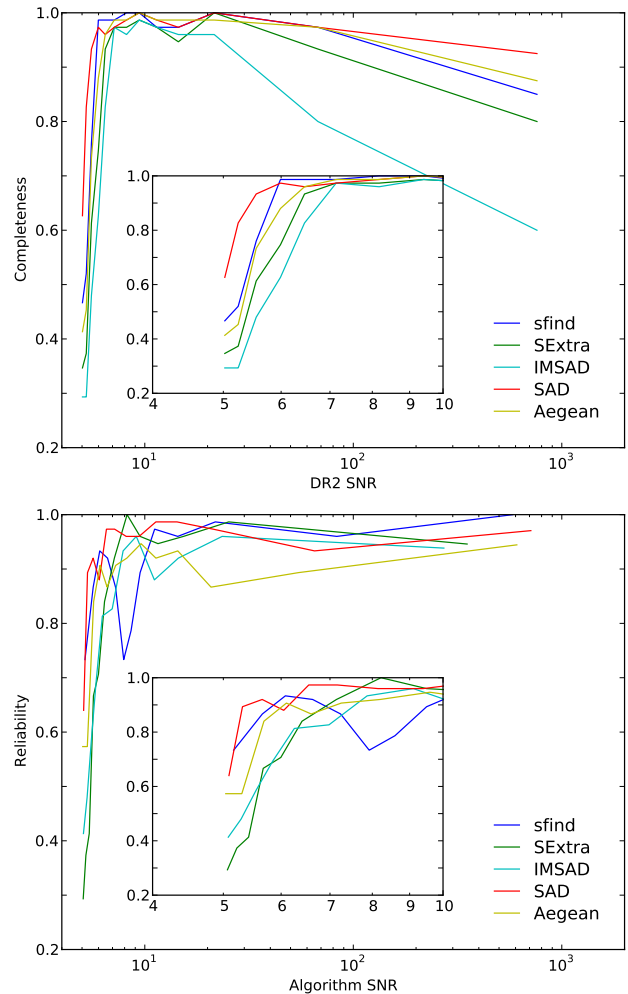

FIG. 9.— The completeness (top) and reliability (bottom) of the catalogs generated by different source-finding algorithms. See §4.1 for the definitions of completeness and reliability used in this work. Results for a background mesh-size (rmsbox) of 20 beamwidths are shown. For `IMSAD`, these results are for the histogram option. For sfind, we have used α=10, whereas for *SExtractor*, `SAD`, `IMSAD` and *Aegean*, detection threshold of 3σ. Only those sources which are $\geqslant 5\sigma$ have then been selected from the respective catalogs prior to comparison with the DR2 catalog. The inset shows only the region where SNR is between 4 and 10. Smoothing over every 75 data points has been done before plotting. Note the low-number statistics for sources with SNR $\gtrsim 70$ as implied by Figure 2.

4.1.4. `IMSAD` *(MIRIAD)*

`IMSAD`[15] computes the image rms noise by fitting a Gaussian to the image histogram, then searches for con-

---

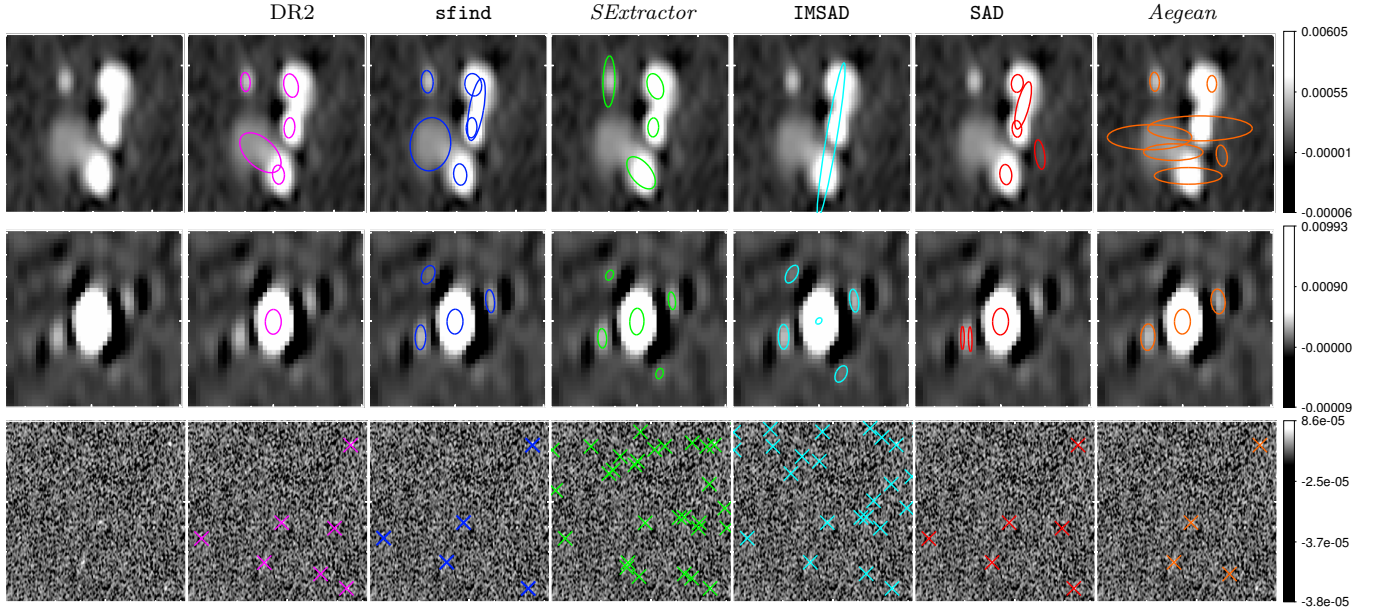[15] www.atnf.csiro.au/computing/software/miriad/doc/imsad.html

FIG. 10.— Examples of 5σ and brighter sources detected by various source-finding algorithms in the DR2 image. Top panel: blended components ($30''$ cutouts centered on $03^h32^m32.2^s, -28°03'09.4''$), middle panel: source with sidelobes ($20''$ cutouts centered on $03^h32^m06.1^s, -27°32'35.8''$), bottom panel: region with a relatively large rms ($12\ \mu$Jy) at the corner of the image ($3'$ cutouts centered on $03^h31^m19.4^s, -27°32'55.6''$). The logarithmic flux density scale shown in each panel has units of Jy. The ellipses have major and minor axes and position angles according to the parameters reported by the respective algorithms.
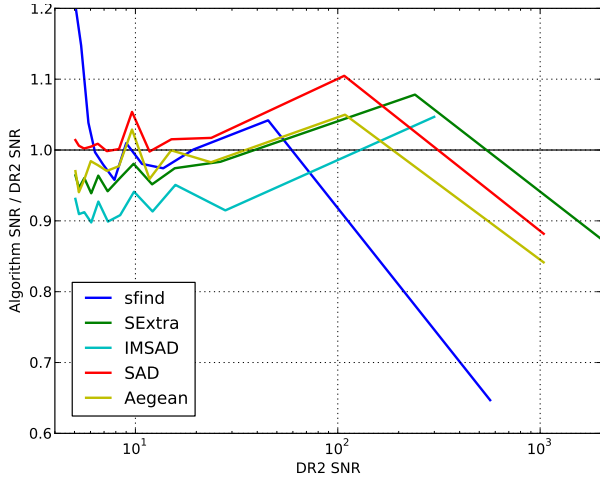

FIG. 11.— The SNR of sources found by different source-finding algorithms relative to their SNR in the DR2 catalog. Smoothing over every 75 data points has been done before plotting. Note the low-number statistics for sources with SNR $\gtrsim 70$ as implied by Figure 2.
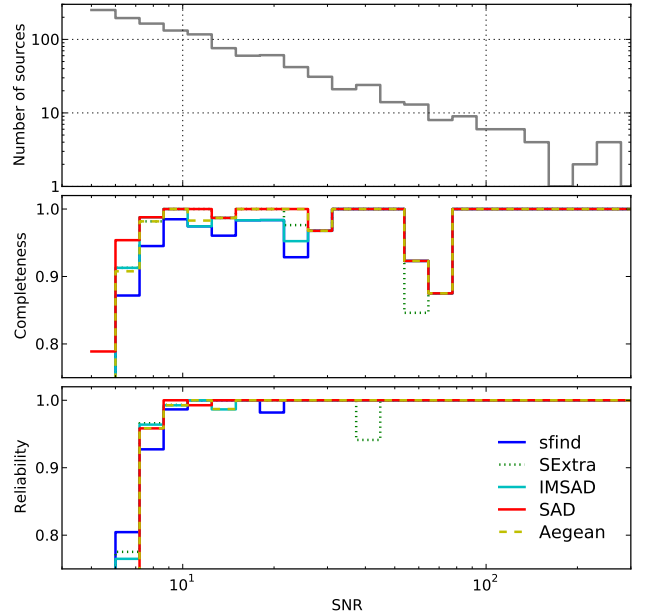

FIG. 12.— Results of the completeness (middle) and reliability (bottom) tests on the Hancock et al. simulated image. Matching radius of $15''$ was used to find counterparts. Input parameters to the source-finding algorithms are same as those given in Figure 9. Here, the completeness and reliability within each SNR bin are plotted (see §4.1), which makes these diagrams different from the ones given in Hancock et al. (2012). For reference, the histogram of the simulated sources is also shown (top). The completeness at SNR~70 dropping below 90% is a result of highly-blended sources.

tiguous pixels (islands) which are above some cutoff and fits the islands with Gaussian components. For fitting, the routine from the *MIRIAD* task `imfit` is used, whereas the island detection is based on the *AIPS* task `SAD`. `IMSAD` can also be used in a mode where the threshold background noise is user-specified. The run-time for this source-finding algorithm searching for sources down to 5σ in the DR2 image is about 2 seconds. We used `IMSAD` version 8-May-98 from MIRIAD 4.2.3 (optimized for CARMA) to carry out these tests.

As before, we searched for sources down to Nσ (specified via the *clip* parameter), and then selected $\geqslant 5\sigma$ sources. The completeness and reliability were tested using the histogram mode (*noplt, hist* options set) as well

as the user-specified mode (only the *noplt* option set; rms of 8 μJy used) for background-rms determination.

The results of these tests on `IMSAD` are given in Table 8. The *noplt, hist* mode usually performs better than the *noplt* mode at least by a few percent in terms of completeness, whereas the latter mode gives better com-

pleteness, again by a few percent. For achieving completeness and reliability of 80% or more, the detection threshold needs to be $\geqslant 6\sigma$. Hence, the optimum use of `IMSAD` would be in this SNR cutoff regime, and the mode in which it is used will depend upon whether completeness or reliability is preferred. Searching down to 3σ followed by rejection of sources below the desired threshold might give better completeness. The reason for the low completeness at a detection threshold of 5σ is that several DR2 sources (which are detected by `sfind`, `SAD` and *Aegean*) close to this threshold are not detected by `IMSAD`, not even at a reduced SNR. The reliability at this threshold is also rather low. This can be attributed to several false point sources detected near the edges of the DR2 image, which are mainly low-SNR features mistaken to be high-SNR due to incorrect assumption of the local rms. Since `IMSAD` does not give the rms for each detected source, one needs to assume this quantity (a constant equal to 8 μJy in our case), and thus, the SNR calculated in regions of increased rms, e.g. regions close to the edge of the DR2 image, is expected to be erroneous. By using an rms image to find the local rms, this false-detection problem can be tackled. Another issue related to false sources is that `IMSAD` does not reject a single sidelobe. All the sidelobes are reported as (genuine) sources. Thus, with respect to false sources detected in noisy regions, missed sources, and sidelobe recognition, `IMSAD` behaves like *SExtractor*. The sources detected by `IMSAD` in some of these scenarios are shown in the image cutouts in Figure 10. Hancock et al. (2012) tested `IMSAD` on a simulated image and found that its completeness and reliability is lower than other source-finding algorithms. However, through our tests on the DR2 image, it appears that `IMSAD` has a competitive performance if used with detection thresholds higher than 5σ.

The completeness of the `IMSAD` catalog as a function of the detection SNR is shown in the lower left panel of Figure 8. The upper panel of Figure 9 shows the `IMSAD` completeness for optimum input parameters. Several sources having 10<SNR<20 as well as SNR$\gtrsim$100, which are detected by other source-finding algorithms, are not detected by `IMSAD`. This is due to the inability of this algorithm to individually fit blended sources and components of extended sources; `IMSAD` tends to fit one elongated source for all components (see upper panel of Figure 10). In Figure 9, the reliability of this algorithm is plotted for the optimum input parameters. `IMSAD` shows reduced reliability between SNRs of 10 and 20, beyond which the reliability roughly flattens off at the 95% level. The reason for this reduced reliability is two-fold. First, several point sources detected close to the edge of the DR2 image have integrated flux densities less than their peak by a factor of a few, clearly indicating unphysical fitting of sources. These false sources can easily be rejected by using a peak-to-integrated flux density ratio criterion. Second, `IMSAD` does not break up islands into components as mentioned above, resulting in extended source positions to be substantially different from those listed in the DR2 catalog. This was also found in the tests carried out by Hancock et al. (2012), which resulted in several false detections. Figure 11 shows the SNR of `IMSAD` sources compared with their counterparts in the DR2 catalog. The assumed constant rms as well as the somewhat lower peak flux density reported by `IMSAD` (for

extended sources) are responsible for the lower SNR with respect to DR2 sources.

4.1.5. *Aegean*

Aegean[16] uses the *FloodFill* algorithm, which separates the foreground pixels from the background and groups them into islands. These "islands" are then passed on to the source characterization stage. Each island of pixels is fit with multiple Gaussian components. The number of components to be fit is determined from a surface-curvature map (cmap), derived from the input image with a Laplacian transform. *Aegean* thus performs a well-constrained multiple Gaussian fitting. A detailed description of *Aegean* and its implementation can be found in Hancock et al. (2012), where the authors demonstrate this source-finding algorithm on a simulated image to produce catalogs for better reliability and completeness than other source finding algorithms. The run-time for *Aegean* searching for sources down to 5σ in the DR2 image is about 4 minutes (using two CPU cores). We used *Aegean* r706 (25-Jul-2012 release) for testing this source-finding algorithm.

We tested the effect of the parameters: (i) *innerclip* (*seedclip*) and *outerclip* (*floodclip*), and (ii) *csigma* on the completeness and reliability of the *Aegean* catalog. innerclip defines the sigma clipping (lower limit) which is used for the detection ("seeding") of islands, whereas outerclip defines that used for reporting ("flooding") the islands in the catalog. *csigma* is the sigma clipping parameter for the curvature map. Note that the r706 version of *Aegean* uses an immutable mesh-size of 20 beamwidths. We used *innerclip* = *outerclip* for our tests. Setting the innerclip to 1σ lower than outerclip gives results alike equating these two clipping parameters, while using a 1σ lower outerclip than innerclip seems to give a few percent better completeness and a slightly reduced reliability.

Our results for *Aegean* are given in Table 9. *csigma* does not have a significant effect on either the completeness or reliability. As with our tests on other source-finding algorithms, we searched for sources down to Nσ. The best completeness and reliability require a 7σ or higher detection threshold. Using a 3σ detection threshold followed by the selection of greater than 7σ might give better completeness at the cost of reliability. Surprisingly, the completeness and reliability are not >95%, as seen in the tests with the Hancock et al. (2012) simulated image. This is because *Aegean* performs well with isolated point sources, but not quite with extended or blended ones. For extended sources, *Aegean* tends to fit some of the components at unexpected locations and with quite elongated Gaussians (the fitting-error flag is also set for these components). Numerous small (but unphysical; size equal to the synthesized beam) components are also reported for several extended sources. Since Hancock et al. (2012) have demonstrated that *Aegean* successfully fits all the components of blended simulated sources, we think that diffuse components and imaging artifacts such as negative sidelobes adversely affect the curvature map produced by this source-finding algorithm, which in turn determines the sources reported. An example of such a case is given in the upper panel

[16] www.physics.usyd.edu.au/%7Ehancock/index.php/Programs/Aegean

of Figure 10. Note that in the latest release of Aegean (r808; 13-Feb-2013) this situation is somewhat improved, since some of the elongated components from r706 are reported with more realistic source sizes (and also without any fitting-error flags), and relatively fewer small-size components are fit.

TABLE 5
COMPLETENESS AND RELIABILITY OF sfind CATALOG

rmsbox	α				
(beamwidths)	10	5	2	1	0.1
Completeness					
5	80.9%	77.8%	74.3%	72.4%	61.9%
10	85.7%	83.7%	79.1%	75.9%	65.5%
20	88.4%	84.5%	81.2%	77.8%	67.8%
50	89.1%	85.7%	82.1%	78.4%	68.0%
Reliability					
5	78.8%	84.9%	90.6%	93.1%	96.4%
10	86.8%	91.6%	95.2%	95.6%	97.4%
20	88.7%	92.5%	95.3%	96.0%	96.4%
50	87.5%	91.4%	93.8%	94.8%	95.3%

Notes— Only $\geqslant 5\sigma$ sources have been selected from the SFIND catalogs prior to comparison with the DR2 catalog.

TABLE 6
COMPLETENESS AND RELIABILITY OF *SExtractor* CATALOG

BACK_SIZE	DETECT_THRESH				
(beamwidths)	3($\geqslant$5)	5	6	7	10
Completeness*					
5	82.2%	72.1%	76.8%	79.1%	77.9%
5 + 3×smoothing	83.4%	73.7%	78.7%	80.0%	79.0%
5 + 10×smoothing	83.3%	73.0%	78.4%	80.0%	79.0%
10	81.7%	71.4%	76.9%	78.7%	78.2%
10 + 3×smoothing	82.1%	71.1%	76.7%	79.1%	77.6%
10 + 10×smoothing	81.7%	71.2%	76.4%	78.6%	77.6%
20	81.6%	70.6%	75.9%	78.6%	77.6%
50	81.3%	70.2%	75.3%	78.0%	77.3%
Reliability					
5	70.4%	87.5%	95.1%	96.3%	97.0%
5 + 3×smoothing	72.4%	89.0%	96.1%	97.2%	97.4%
5 + 10×smoothing	72.2%	89.5%	96.1%	97.2%	97.4%
10	75.2%	90.8%	96.2%	96.5%	97.4%
10 + 3×smoothing	75.8%	90.1%	96.2%	97.0%	97.4%
10 + 10×smoothing	75.6%	90.2%	96.2%	96.9%	97.4%
20	77.0%	90.3%	96.4%	96.9%	97.4%
50	78.1%	91.1%	96.5%	96.9%	97.4%

Figure 8 shows the completeness of the *Aegean* catalog as a function of SNR of sources in the DR2 catalog, and the upper panel of Figure 9 shows the completeness for $csigma=1\sigma_{\mathrm{cmap}}$, and searching for sources down to 3σ. *Aegean* completeness appears to be quite good except for sources below 7σ in the DR2 catalog, which are detected at a decreased SNR (below 5) and hence rejected from the catalog, and for a few components of extended sources above $\sim 100\sigma$, which are reported to have positions not matching those in the DR2 catalog. The decreased SNR reported for 5–7σ sources is due to a slightly decreased peak flux density and a slightly increased rms with respect to the DR2 catalog. Figure 9 plots the reliability using $csigma=1\sigma_{\mathrm{cmap}}$, and searching for sources down to 3σ. The significant deviation of reliability from unity for SNR>10 sources, results from the several (false) small-

TABLE 7
COMPLETENESS AND RELIABILITY OF SAD CATALOG

IMSIZE	CPARM				
(beamwidths)	3($\geqslant$5)	5	6	7	10
Completeness*					
5	77.8%	75.5%	80.0%	83.7%	78.4%
10	89.9%	88.7%	93.1%	92.1%	93.4%
20	93.4%	92.2%	95.9%	95.4%	97.1%
50	95.9%	94.8%	96.5%	95.8%	97.1%
actnoise‡	87.4%	86.7%	95.2%	96.7%	95.7%
Reliability					
5	88.1%	89.0%	97.3%	99.1%	99.3%
10	91.2%	93.0%	97.4%	97.4%	96.6%
20	92.5%	93.9%	96.7%	96.9%	97.0%
50	85.1%	87.3%	88.1%	88.6%	92.1%
actnoise‡	58.2%	60.2%	86.6%	94.6%	96.4%

Notes— ‡Search using the actnoise keyword in the FITS header (=7.465E-06 JY/BM).

TABLE 8
COMPLETENESS AND RELIABILITY OF IMSAD CATALOG

options	clip				
	3($\geqslant$5)	5	6	7	10
Completeness*					
noplt‡	75.7%	73.9%	78.1%	84.2%	79.3%
noplt,hist	75.4%	76.5%	91.1%	90.5%	90.6%
Reliability					
noplt‡	77.4%	83.7%	95.3%	96.5%	96.8%
noplt,hist	78.2%	79.2%	91.7%	95.8%	96.2%

Notes— ‡Clipping level is manually entered as the appropriate multiple of the background rms chosen to be 8 μJy.

TABLE 9
COMPLETENESS AND RELIABILITY OF *Aegean* CATALOG

csigma	innerclip = outerclip				
(cmap rms)	3($\geqslant$5)	5	6	7	10
Completeness*					
0.5	86.1%	78.7%	83.8%	86.1%	85.8%
1.0	86.3%	78.5%	84.2%	86.4%	85.8%
2.0	88.2%	78.5%	84.5%	86.8%	85.5%
Reliability					
0.5	86.2%	91.6%	91.6%	92.1%	91.7%
1.0	84.9%	90.6%	90.5%	90.5%	91.1%
2.0	83.4%	91.0%	91.2%	91.4%	89.9%

Notes— *The fraction of sources in the DR2 catalog which are $\geqslant 6\sigma$, 7σ and 10σ are 73.0%, 58.1% and 37.7% respectively. Completeness for these detection thresholds has been normalized accordingly. A detection threshold of 3($\geqslant$5) implies a search down to 3σ followed by the selection of only those source that are greater than 5σ.

size, and elongated components reported for extended sources, as mentioned above. Figure 11 shows the SNR of sources detected by *Aegean*, in comparison with the corresponding sources in the DR2 catalog. The disagreement in SNR for 5–7σ sources in the DR2 catalog has been explained above. For SNR>200 sources, the peak flux densities agree quite well between the Aegean and DR2 catalogs, but the rms reported by Aegean is consistently higher, which causes the disagreement at the high-SNR end of the diagram.

4.1.6. *Summary of results from the efficiency tests*

We find remarkable differences between algorithms in terms of components fitted for extended sources, sidelobe rejection, and point sources detected in regions where the rms is appreciably larger than the mean rms. Our results for completeness and reliability are broadly similar to those of Huynh et al. (2011) and Hancock et al. (2012). For applications that need both completeness as well as reliability, sfind and *Aegean* are good. Additionally, we found that the SAD algorithm within the widely available *AIPS* package had a better performance. IMSAD also gives a good completeness and reliability for detection thresholds $\geqslant 6\sigma$. For transient searches, reliability takes preference over completeness, since false positives are likely to consume follow-up resources. Most transient projects are likely to be searching in near real-time. However, in this particular project we were fortunate to have a deep reference image that was more than three times deeper than the single-epoch images. This allowed us to study reliability with real (rather than synthetic) datasets. From Figure 9, we see that the reliability of sfind is better than that of *SExtractor*, IMSAD and *Aegean* except for SNR near 8. However, the best reliability is provided by SAD.

4.2. *Transient Candidate Search*

Using the best-performing source finding algorithms from §4.1.6 and their optimum parameter values, we carried out a search for transient radio sources over all epochs. We ran sfind, SAD, IMSAD and *Aegean* on single-epoch images and obtained 49 single-epoch catalogs for each algorithm. We required that any potential transient candidate identified in the single-epoch catalog obey the following constraints.

1. The source is not found in the reference catalog (within 2″) of persistent sources. The reference catalog was constructed similar to DR2 but the sources were selected up to the 20% power point of the beam (i.e. $\theta < 21.5'$ radius) of each pointing rather than the 34′ interior region shown in Figure 1.

2. The source is at least a 7σ detection.

3. Is a genuine point-like source, i.e. it has

   (a) $0.9 < S_{int}/S_{peak} < 1.5$
   (b) $a < 2.8'' \times 1.5$, $b < 1.6'' \times 1.5$
       (a =major axis, b =minor axis; recall that the synthesized beam is $2.8'' \times 1.6''$)

4. The source is at least 20 synthesized beams (20 × geometric mean of FWHMs; 42″) away from the nearest:

   (a) bright source (> 500 μJy), so that any sidelobe emission is rejected
   (b) extended source

The multiplicative factor of three-halves used in the selection of point sources, as well as the distance of 20 beamwidths used for constraining the proximity from bright and extended sources, is somewhat arbitrary, but is based on several iterations of our transient search code and inspection of the cutouts of the resulting transient candidates. By investigating how the major axes of sources increase with their distance from the pointing center, we found that $\theta \simeq 21.5'$ appears also to be the threshold beyond which bandwidth smearing coupled with our constraints on the major (and minor) axes start rejecting genuine point sources.

Due to the large number of synthesized beams searched ($n = 1.8 \times 10^7$) in this dataset, there is a modest probability that a transient candidate is due to noise[17]. We thus carried out an analysis similar to Frail et al. (2012) (see Appendix A of that paper) to determine the SNR above which the probability of having the highest value of n Gaussian random numbers is $\leqslant 1\%$. This corresponds to an SNR of 6.1. However, following the recommendation of Frail et al. (2012) to have a higher SNR cutoff when the noise was not strictly Gaussian, we chose 7σ as the lower limit for finding transients.

The search method outlined above may miss transients which are bright enough to be present in the reference catalog. Therefore, we also searched (with similar constraints as above) for sources which are detected in the reference catalog of persistent sources and detected in only one of the single-epoch catalogs.

For sfind, we adopted the parameters α=10 and $rmsbox$=20 for transient search. We found five candidates which are present only in a single epoch above the 7σ detection threshold. They also show up in the reference catalog. However, all of these candidates seem to be variable but persistent, and are detected at a low SNR ($\sim$3–5σ) in other epochs. Thus we do not find any transient with sfind.

For *Aegean*, we used a clipping level of 7σ and a curvature-map cutoff of 1σ. We discarded all the transients corresponding to islands which were too small to give a 6-parameter Gaussian fit (sources with flag 10000 or 00100 set) since otherwise we were dealing with a large number of transients. Using these constraints, we found one candidate which is present in the reference catalog and is a persistent source detected at a largely reduced SNR in other epochs. Accordingly, *Aegean* does not yield any transients.

For SAD, we searched for sources down to 3σ with an rmsbox of 20 beamwidths, and used the input parameters $DPARM(3), DPARM(7)$=(1000,1) to ensure optimum completeness and reliability based on our efficiency test. We do not find any transients with SAD.

With IMSAD, we searched for sources down to 7σ with the histogram option set. The several sources found to be transient candidates, are only variables at our detection limit, as evident through visual inspection of the images. Thus, there are no transients reported by IMSAD.

5. DISCUSSION AND CONCLUSIONS

In this paper we have explored the time-domain properties of a 1.4 GHz survey made toward the Extended Chandra Deep Field South (E-CDFS) region. Six mosaic pointings toward E-CDFS were taken in 49 separate epochs over a period of three months. Single-epoch images allow us to explore the transient and variable radio sky at sub-milliJansky levels on timescales of days, weeks

---

[17] From theory we know that the statistics of beam values of interferometric maps should follow a Gaussian distribution.

and months. We will now use these data to assess the degree of variability (§5.1) and the transient rate (§5.2) of the radio sky, and predict what will be seen by future wide-field surveys (§5.3).

5.1. *Comparison of variability with previous surveys*

We found in our study that only a small fraction ($7/599 = 1.2^{+1.2}_{-0.7}\%$) of the point sources in the E-CDFS showed any significant variability on day-week-month timescales. Evidently, the sub-milliJansky radio sky at 1.4 GHz is not highly variable. The only previous sub-mJy study at 1.4 GHz was from a single deep pointing toward the Lockman Hole. With sampling timescales of 19 days, Carilli et al. (2003) found less than 2% of sources above 0.1 mJy to be highly variable.

Our findings at sub-mJy levels are consistent with the several previous 1.4 GHz studies at higher flux density thresholds. Thyagarajan et al. (2011) analyzed the 8444 deg^2 of the FIRST survey and found only 0.5% (1627/279407) of sources above 1 mJy varied significantly on timescales of minutes to years. Frail et al. (1994) imaged a 2° region toward a gamma-ray burst on timescales of 1-96 days and found that fewer than ∼1% of the sources above a flux density of 3.5 mJy were strongly variable. de Vries et al. (2004) imaged a 120.2 deg^2 area of Sloan Digital Sky Survey Stripe 82 finding ⩽1.4% (123/9086) of the radio sources to be strongly variable (i.e. >4σ) above flux densities of 2 mJy on a 7 year timescale. Ofek & Frail (2011) do a two-epoch comparison of FIRST and NVSS point sources brighter than 5 mJy and find that only 0.1% (43/4367) vary by more than 4σ over timescales ranging from about 300 to 1700 days. Bannister et al. (2011a,b) analyzed 22 years and 2775 deg^2 of MOST observations at 0.84 GHz and found only 0.17% (53/29730) strong variables above 14 mJy on timescales of days to minutes to years. Finally, Croft et al. (2011) used the Allen Telescope Array to survey a 690 deg^2 area at 1.4 GHz. They compared their catalog to the NVSS, finding that 0.1% (6/4408) of the sources were highly variable on a timescale of 15 years.

In Figure 13 we plot the differential source counts for the persistent radio sky at 1.4 GHz, normalized in the usual way by the Euclidean rate (Huynh et al. 2005). The steep evolution of the AGN with decreasing flux density is apparent, as is the flattening of the source counts near 1 mJy. The fractional variability appears to be low, at a level of one percent, among the sources greater than 100 μJy.

Comparing the variability properties at 1.4 GHz of our sub-mJy population with those at higher flux densities, we find no obvious change in the *fraction* of strong variables. This is despite the fact that radio-loud AGN (which dominate at milliJansky levels) diminish in importance as radio-quiet AGN and star forming galaxies begin to populate the radio sky at lower flux densities, the latter constituting about 50% of the sub-milliJansky radio sky (Smolčić et al. 2008; Padovani 2011). On the other hand, in a sample of brightness temperature-limited radio sources, it would be expected that the variability would increase with decreasing flux density.

Our optical identifications for the seven radio variables show a mix of AGN and star-forming galaxies. Six of these are consistent with a nuclear source based on the carefully matched astrometry. There is no evidence from this sub-mJy sample that we are seeing a new source of variable radio emission such as supernovae, gamma-ray bursts, etc. that would be expected to be offset from the nucleus.

Radio variability appears to be a strong function of frequency. The best study to compare to this one is the 5 GHz survey of Ofek et al. (2011), since it used a similar observing cadence and the identical statistical measures of variability. In this case it was found that ∼30% of point sources brighter than 1.5 mJy at 5 GHz were variable.

Some caution is warranted in comparing fractional variability between different experiments. Differences in the cadence, integration time, duration, and angular resolution will have a tendency of reducing the variability amplitude. For example, the low angular resolution of some surveys may have the effect of reducing the strong variability from compact radio sources embedded in diffuse emission. This current survey with a synthesized beam of $2.8'' \times 1.6''$ has the highest angular resolution of any previous variability survey (see Ofek et al. 2011).

Another factor to consider is that the degree of variability which can be detected in a given experiment, as measured by the modulation index m, is a function of the signal-to-noise ratio. A source in our present survey would have to have a mean flux density of 1 mJy in order to detect a 10% modulation, while at the 5σ limit of the DR2 catalog (∼40 μJy) a source would have to have $m \geqslant 1.5$ in order to be identified as a significant variable. No strong variables (i.e. $m > 0.5$) were identified in our survey, but only the about 90 sources in the DR2 catalog are bright enough ($\geqslant 126$ μJy) to have been identified as a strong variable. In either case the fraction of significant or strong variables is less than a few percent of the sample.

The most robust conclusion that can be drawn is that the variable radio sky at 1.4 GHz appears to be relatively quiet with only a fraction of a percent of sources varying substantially over a wide range of flux densities and timescales.

5.2. *Limits on transient areal density and rate*

We searched our multi-epoch data for transients but found none. The search was conducted on each image out to a radius of 21.5' from the pointing center. The single-epoch area out to that radius is 0.40 deg^2, or a total area of 20 deg^2 for all 49 epochs. However, the sensitivity of the VLA antennas is not uniform across this area. The primary beam response is well-described by a Gaussian with a half-width to half-maximum of 15', falling to 20% response at our search radius of 21.5'. At the pointing center the 7σ flux density limit was approximately 210 μJy for each epoch.

In order to calculate a limit on the areal density of any putative transient population we follow Ofek et al. (2011) and parameterize the source number-count function as a power law of the form

$$\kappa(> S) = \kappa_0 (S/S_0)^{-\gamma}, \qquad (4)$$

where S is the peak flux density, $\kappa(> S)$ is the sky surface density of sources brighter than S, κ_0 is the sky surface density of sources brighter than S_0, and γ is the power
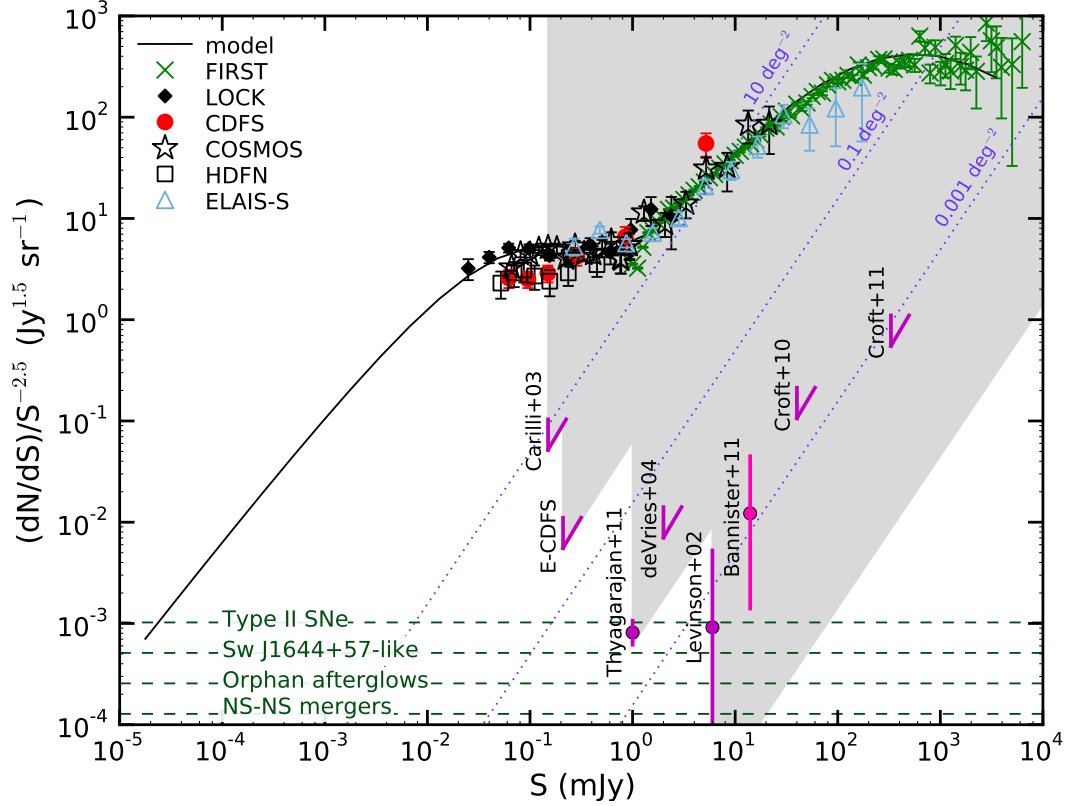
FIG. 13.— Normalized 1.4 GHz differential radio source counts for persistent sources from de Zotti et al. (2010) and the normalized areal density of transients (or limits) as a function of the flux density for various surveys at this frequency. Bannister et al. (2011a) survey at 0.84 GHz is colored differently than the other surveys. Most of the surveys are upper limits (wedge symbols) and the sampled phase space is shown by the grey shaded area. Upper limits from Frail et al. (1994) and Bower & Saul (2011) do not explore any new part of the phase space (non-grey area), and hence have been left out of this diagram. Our upper-limit is labeled as "E-CDFS". Three surveys have transient detections so far, 2σ error bars for which are shown according to Gehrels (1986). Note that Thyagarajan et al. (2011) and Bannister et al. (2011a) may have identified a few strong variables as transients (see §5.2), which would make their detections move downward on this plot. The black solid line is the model for AGNs and star-forming galaxies from Condon (1984). Lines of constant areal density are shown as blue dotted lines. The horizontal dashed lines are estimates for the areal density for known and expected classes of long-duration radio transients taken directly from Frail et al. (2012). The areal density for Swift J1644+57-like tidal disruption events has been modified according to Berger et al. (2012) to reflect their true rate at 1.4 GHz. Upper limits from the ASKAP-VAST surveys are estimated to be an order of magnitude or more below the rate of orphan gamma-ray burst afterglows, and having a rms sensitivity ranging between 10 μJy and 0.5 mJy.

law index of the source number-count function. We assume, for simplicity, a homogeneous source distribution in Euclidean Universe so $\gamma = 3/2$. The one-sided 2-σ upper limit on the areal density is three events (Gehrels 1986). Therefore, using Equation C5 in Ofek et al. (2011) we find that the 2σ upper limit on areal density to a flux limit of 210 μJy is 18.0 deg^{-2} per epoch. Given that we have 49 epochs the 2σ upper limit on the areal density is $\kappa(> 0.21\,\mathrm{mJy}) < 0.37\,\mathrm{deg}^{-2}$. We can further estimate an upper limit on the transient rate assuming a duration t_{dur} less than the shortest time between epochs of $\Re(> 0.21\,\mathrm{mJy}) < 268(t_{\mathrm{dur}}/0.5\,\mathrm{day})^{-1}\ \mathrm{deg}^{-2}\,\mathrm{yr}^{-1}$.

Our upper limit on the areal density of transient sources at sub-milliJansky levels can be compared with the predictions based on previous surveys. The Bower et al. (2007) survey is a useful benchmark since their areal density dominates all known classes of transients. Adopting their measured two epoch rate of $\kappa(> 0.37\,\mathrm{mJy}) = 1.5\,\mathrm{deg}^{-2}$ and assuming a Euclidean source distribution (i.e. $\gamma = 3/2$) we predict $\kappa(> 0.21\,\mathrm{mJy}) = 3.5\,\mathrm{deg}^{-2}$ at the flux density limit of our current survey.

An alternative way to look at our results is to compare our null detection to the expected number of Bower et

al. transients expected in our dataset. We use the parameterization of Fender & Bell (2011) for the predicted Bower et al. transient rate as a function of flux density,

$$\log\left(\frac{\kappa}{\mathrm{deg}^{-2}}\right) = -1.5\,\log\left(\frac{S_\nu}{\mathrm{Jy}}\right) - 5.13 \qquad (5)$$

where κ is the snapshot rate, and S_ν denotes the detection threshold of the observations at the pointing center (i.e. $7\sigma = 210\ \mu$Jy). Integrating both sides of equation 5 over the azimuthal angle and in θ out to 21.5′ we get about 0.42 transients per epoch, if the Bower et al. (2007) transients are real. Since we have 49 epochs, we expect to have about 21 Bower et al. transients in our E-CDFS dataset.

Our search on the E-CDFS field suggests that the areal density of radio transients is an order of magnitude or more below the rate measured by Bower et al. (2007) (i.e. $<0.37\ \mathrm{deg}^{-2}$ vs. $3.5\,\mathrm{deg}^{-2}$). Alternatively, we find a 2σ upper limit of <3 transients while the predicted number is $\sim$21 transients. Our work therefore appears to support that of Frail et al. (2012), which found from a re-analysis of the Bower et al. data, that the transient rate was as

much as an order of magnitude smaller than previously reported. This conclusion would be more robust if the spectral index of the putative transient population was better known. Ofek et al. (2010) was able to use other surveys to constrain the spectra index $\alpha > 0$ (where $S_\nu \propto \nu^\alpha$). Since the Bower et al. (2007) rates were derived based on observations made mostly at 5 GHz, our only data provide strong constraints for $0 \leqslant \alpha \leqslant 1.1$. A population of optically thick $\alpha \simeq 5/2$ sources with a rate similar to that of Bower et al. (2007) would be undetected in our 1.4 GHz E-CDFS fields.

Many of the same variability surveys discussed in §5.1 were also sensitive to transients. Superposed on the radio source-count plot of Figure 13 are the results of several of these transients surveys. Light grey shaded areas represent the transient phase space covered by each of the surveys, white space represents open phase space for future narrow-deep or wide-shallow surveys. With few exceptions most of these transient surveys result in upper limits. Also shown are the normalized areal density of several known and expected classes of long-duration radio transients, based on Frail et al. (2012). The nominal rates for the putative Bower et al. (2007) sources are about two orders of magnitude above the tidal disruption or Sw J1644+57-like objects. Thyagarajan et al. (2011) define a transient as having either a single detection in the analyzed epochs, or the highest flux density 5 times greater than the next highest one (detection/upper-limit). Thus, Thyagarajan et al. (2011) and 2-epoch surveys like Bannister et al. (2011a) may identify a strong variable as a true transient, which will move the source-count of detected transients lower in Figure 13.

5.3. *Future Radio Surveys*

There are several facilities built or under construction that will be capable of synoptic imaging at 1.4 GHz. All of these facilities have the exploration of the time domain as part of their core science programs. The Australian Square Kilometer Array Pathfinder (ASKAP) and the Apertif instrument on the Westerbork Synthesis radio Telescope (WSRT) will be using focal plane array technology to instantaneously image an instantaneous field of view (FoV) of 30 deg^2 and 8 deg^2, respectively (Murphy et al. 2012; Oosterloo et al. 2010). South Africa is building MeerKAT, an array of 64 13.5-m diameter dishes, with a FoV of 1 deg^2 at 1.4 GHz (Booth et al. 2009). Finally, there is the newly refurbished Karl G. Jansky Array (VLA) has 27 25-m dishes with a FoV of 0.25 deg^2 at 1.4 GHz (Perley et al. 2011).

We list the capabilities of each of these telescopes in Table 10. Survey speed (SS), normalized here to the VLA, is a useful figure of merit for inter-comparison of survey capabilities of long duration transients and is expressed as

$$SS \propto BW \times \Omega(A_e/T_{sys})^2 \qquad (6)$$

where BW is the bandwidth, Ω the FoV, A_e is the total collecting areas times the aperture efficiency ϵ_e, and T_{sys} is the antenna system temperature (Cordes 2008). The relative survey speeds are only approximate since some of the system parameters for MeerKAT and ASKAP have not been confirmed with on-the-sky testing. Like-

wise, we have assumed that radio frequency interference (RFI) limits the VLA bandwidth to only 50% of its maximum BW. Larger fractions are achievable in the more extended array configurations and with better RFI excision of the data. We have calculated the ASKAP survey speed with both the 18 phased array feeds which are currently funded and with the full 36 as originally specified. Despite these uncertainties, it is clear from Table 10 that to within factors of a few, these are all powerful wide-field imaging facilities.

TABLE 10
TELESCOPE SPECIFICATIONS

Telescope	BW (MHz)	Ω (deg^2)	D (m)	N	ϵ_e	T_{sys} (K)	SS
VLA	512	0.25	25	27	0.5	26	1.0
ASKAP	300	30	12	18(36)	0.8	50	1.1(4.6)
Apertif	300	8	25	13	0.75	70	1.3
MeerKAT	750	1	13.5	64	0.7	30	4.1

Notes— Here BW is bandwidth in MHz, Ω is the field of view in deg^2, D is the antenna diameter in meters, N is the number of antennas in the array, ϵ_e is the aperture efficiency, T_{sys} is the system temperature in Kelvin, and SS is the survey speed normalized to the VLA.

To illustrate these survey capabilities and compare them to what we currently know about the transient and variable radio sky, we will use the example of a electromagnetic counterpart search for gravitational waves. For a good overview of the topic of EM-GW searches, and the main issues, we refer the reader to Metzger & Berger (2012) and Nissanke et al. (2011).

Long-duration radio emission has been predicted to originate from the merger of a neutron star binary from several sources including the merger shock (Kyutoku et al. 2012), afterglow emission from the beamed outflow in the relativistic and non-relativistic phases (Metzger & Berger 2012), and from quasi-isotropic, mildly relativistic outflows ejected during the merger (Nakar & Piran 2011; Piran et al. 2012). All of these mechanisms depend on the amount of energy put into shocked material and the density of the ambient medium. Predicted flux densities and timescales therefore vary over a wide range. We take as an example the detection of a signal with a peak flux density of 100 μJy. Such a signal might be expected to occur for an ambient medium with a density of 0.1 cm^{-3} on day timescales for a merger shock, or on year-long timescales for mildly relativistic ejecta. We note that if short duration gamma-ray bursts are the dominant population of neutron star mergers, then the predicted radio signal would be difficult to detect with the telescopes in Table 10, given the the canonical energy and ambient density inferred for this population (Fong et al. 2012; Metzger & Berger 2012).

The median sky localization of a gravitational-wave source will be 60 deg^2 with a three-element GW network, and 7 deg^2 with a five- element GW network (Nissanke et al. 2011). With such large error boxes, the main challenge for the identification of a EM-GW counterpart will be to distinguish it from the foreground of false positives. Optical-only searches for EM-GW counterparts are expected to be overwhelmed by false positives at the required depth of 22-23 mag and special strategies are required (Nissanke et al. 2011). However, as we have shown

from this paper, the radio sky at 1.4 GHz is relatively quiet. Integrating the differential source counts in Figure 13 using the fit from Huynh et al. (2005), we estimate that the number of *persistent* radio sources above 100 μJy to be 910 deg^{-2}. We have estimated that the fraction of strong variables, on a wide range of timescales, is likely to be one percent or less, or 9 strong variables per square degree. Further, the radio variables that we have seen to date have all been nuclear sources (§3.3). Such variable or transient sources could be easily rejected as EM-GW counterparts since significant offsets are predicted from the host galaxy based on binary neutron-star population synthesis models and measurements from short-hard gamma-ray bursts (Belczynski et al. 2006; Fong et al. 2010). Lastly, the number of unrelated transients is also expected to be much less. Our derived limit on the transient areal rate of $\kappa(> 0.21 \, \mathrm{mJy}) < 0.37 \, \mathrm{deg}^{-2}$ translates to a $<1.1 \, \mathrm{deg}^{-2}$ at 100 μJy (§5.2) for a Euclidean distribution. The limit on the *known* transient populations is even smaller (Frail et al. 2012).

We believe that the multi-wavelength approach which we have taken here should inform future searches. For the radio variables that we found in the E-CDFS, we were able to identify the source of the emission using optical images and spectra (§3.3). A similar strategy could be employed to identify false-positives for the small number of radio variables or transient sources identified in EM-GW counterpart searches. Whether this approach will ultimately lead to a robust EM-GW counterpart detection is uncertain, but in terms of characterizing the variables, minimizing false-positives and getting an early sense on the nature of transients, we suggest that joint radio-optical searches will be fruitful for exploring the dynamic sky.

REFERENCES

Afonso, J., Mobasher, B., Koekemoer, A., Norris, R. P., Cram, L. 2006, AJ, 131, 1216
Assef, R. J., Kochanek, C. S., Brodwin, M., et al. 2010, ApJ, 713, 970
Bannister, K. W., Murphy, T., Gaensler, B. M., Hunstead, R. W., & Chatterjee, S. 2011, MNRAS, 412, 634
Bannister, K. W., Murphy, T., Gaensler, B. M., Hunstead, R. W., & Chatterjee, S. 2011, Erratum, July 2011
Belczynski, K., Perna, R., Bulik, T., et al. 2006, ApJ, 648, 1110
Bell, M. E., Fender, R. P., Swinbank, J., et al. 2011, MNRAS, 415, 2
Berger, E., et al. 2001, Nature, 410, 338
Berger, E., Zauderer, A., Pooley, G. G. et al. 2012, ApJ, 748, 36
Bonzini, M., Mainieri, V., Padovani, P., et al. 2012, ApJS, 203, 15
Booth, R. S., de Blok, W. J. G., Jonas, J. L., & Fanaroff, B. 2009, arXiv:0910.2935
Bower, G. C., & Saul, D. 2011, ApJ, 728, L14
Bower, G. C., Saul, D., Bloom, J. S., Bolatto, A., Filippenko, A. V., Foley, R. J., & Perley, D. 2007, ApJ, 666, 346
Bridle, A. H., & Schwab, F. R. 1989, ASPC, 6, 247
Burke-Spolaor, S., Bailes, M., Ekers, R., Macquart, J.-P., & Crawford, F., III 2011, ApJ, 727, 18
Burrows, D. N., Kennea, J. A., Ghisellini, G., et al. 2011, Nature, 476, 421
Cameron, P. B., et al. 2005, Nature, 434, 1112
Camilo, F., Ransom, S. M., Halpern, J. P., Reynolds, J., Helfand, D. J., Zimmerman, N., & Sarkissian, J. 2006, Nature, 442, 892
Carilli, C. L., Ivison, R. J., & Frail, D. A. 2003, ApJ, 590, 192
Cenko, S. B., Krimm, H. A., Horesh, A., et al. 2012, ApJ, 753, 77
Chary, R., & Elbaz, D. 2001, ApJ, 556, 552
Condon, J. J., Cotton, W. D., Fomalont, E. B., et al. 2012, ApJ, 758, 23
Condon, J. J., 2007, Astronomical Society of the Pacific Conference Series, 380, 189
Condon, J. J., Cotton, W. D., Broderick, J. J., 2002, AJ, 124, 675
Condon, J. J., 1984, ApJ, 284, 44
Cordes, J. M. 2008, Frontiers of Astrophysics: A Celebration of NRAO's 50th Anniversary, 395, 225
Cotton, W. D. & Perley, R. 2010, OBIT Development Memo Series No. 17, ftp://ftp.cv.nrao.edu/NRAO-staff/bcotton/Obit/EVLABeam.pdf
Cotton, W. D. & Peters, W. 2011, OBIT Development Memo Series No. 25, ftp://ftp.cv.nrao.edu/NRAO-staff/bcotton/Obit/FDR.pdf
Croft, S., Bower, G. C., Keating, G., Law, C., Whysong, D., Williams, P. K. G., & Wright, M. 2011, ApJ, 731, 34
Damen, M., Labbé, I., van Dokkum, P. G., et al. 2011, ApJ, 727, 1
Desai, V., Armus, L., Spoon, H. W. W., et al. 2007, ApJ, 669, 810
Donley, J. L., Koekemoer, A. M., Brusa, M. et al. 2012, ApJ, 748, 142
Fender, R. P., & Bell, M. E. 2011, Bulletin of the Astronomical Society of India, 39, 315

Fong, W., Berger, E., Margutti, R., et al. 2012, ApJ, 756, 189
Fong, W., Berger, E., & Fox, D. B. 2010, ApJ, 708, 9
Fox, D. B., et al. 2005, Nature, 437, 845
Frail, D. A., Kulkarni, S. R., Ofek, E. O., Bower, G. C., & Nakar, E. 2012, ApJ, 747, 70
Frail, D. A., Kulkarni, S. R., Berger, E., & Wieringa, M. H. 2003, AJ, 125, 2299
Frail, D. A., Kulkarni, S. R., Hurley, K. C., Fishman, G. et al. 1994, ApJ, 437, 43
Gaensler, B. M., et al. 2005, Nature, 434, 1104
Gaensler, B. M., & Hunstead, R. W. 2000, PASA, 17, 72
Gehrels, N. 1986, ApJ, 303, 336
Giavalisco, M., Ferguson, H. C., Koekemoer, A. M., et al. 2004, ApJ, 600, 93
Hancock, P. J, Murphy, T., Gaensler, B. M, Hopkins, A., & Curran, J. R 2012, MNRAS, 422, 1812
Hallinan, G., et al. 2007, ApJ, 663, L25
Hankins, T. H., Kern, J. S., Weatherall, J. C., & Eilek, J. A. 2003, Nature, 422, 141
Hopkins, A. M., Miller, C. J., Connolly, A. J., et al. 2002, AJ, 123, 1086
Huynh, M. T., Jackson, C. A., Norris, R. P., & Prandoni, I. 2005, AJ, 130, 1373
Huynh, M., Hopkins, A., Norris, R., et al. 2012, PASA, 29, 229
Huynh, M. T., Hopkins, A. M., Lenc, E., et al. 2012, arXiv:1208.6306
Hyman, S. D., Lazio, T. J. W., Kassim, N. E., Ray, P. S., Markwardt, C. B., & Yusef-Zadeh, F. 2005, Nature, 434, 50
Hyman, S. D., Wijnands, R., Lazio, T. J. W., Pal, S., Starling, R., Kassim, N. E., & Ray, P. S. 2009, ApJ, 696, 280
Ibar, E., Ivison, R. J., Biggs, A. D., et al. 2009, MNRAS, 397, 281
Johnston, S., Taylor, R., Bailes, M., et al. 2008, Experimental Astronomy, 22, 151
Kellermann, K. I., Fomalont, E. B., Mainieri, V., et al. 2008, ApJS, 179, 71
Komatsu, E., Smith, K. M., Dunkley, J. et al. 2011, ApJS, 192, 18K
Körding, E., Rupen, M., Knigge, C., Fender, R., Dhawan, V., Templeton, M., & Muxlow, T. 2008, Science, 320, 1318
Kyutoku, K., Ioka, K., & Shibata, M. 2012, arXiv:1209.5747
Lazio, J., Bloom, J. S., Bower, G. C., Cordes, J., Croft, S., Hyman, S., Law, C., & McLaughlin, M. 2009, astro2010: The Astronomy and Astrophysics Decadal Survey, 2010, 176
Levinson, A., Ofek, E. O., Waxman, E., & Gal-Yam, A. 2002, ApJ, 576, 923
Lovell, J. E. J., et al. 2008, ApJ, 689, 108
Magnelli, B., Elbaz, D., Chary, R. R., et al. 2011, å, 528, 35
Magnelli, B., Elbaz, D., Chary, R. R., et al. 2009, å, 496, 57
Mainieri, V., Kellermann, K. I., Fomalont, E. B., Miller, N., et al. 2008, ApJS, 179, 195
McLaughlin, M. A., et al. 2006, Nature, 439, 817
Metzger, B. D., & Berger, E. 2012, ApJ, 746, 48

Miller, N. A., Fomalont, E. B., Kellermann, K. I., Mainieri, V., Norman, C., Padovani, P., Rosati, P., & Tozzi, P. 2008, ApJS, 179, 114

Miller, N. A., Bonzini, M., Fomalont, E. B., et al. 2013, ApJS, 205, 13

Murphy, T., Chatterjee, S., Kaplan, D. L., et al. 2012, arXiv:1207.1528

Nakar, E., & Piran, T. 2011, Nature, 478, 82

Nissanke, S., Kasliwal, M., Georgieva, A. 2013, ApJ, submitted

Norris, R., P., Afonso, J., Appleton, P. N., et al. 2006, AJ, 132, 2409

Ofek, E. O., Breslauer, B., Gal-Yam, A., Frail, D., Kasliwal, M. M., Kulkarni, S. R., & Waxman, E. 2010, ApJ, 711, 517

Ofek, E. O., & Frail, D. A. 2011, ApJ, 737, 45

Ofek, E. O., Frail, D. A., Breslauer, B., et al. 2011, ApJ, 740, 65

Oosterloo, T., Verheijen, M., & van Cappellen, W. 2010, ISKAF2010 Science Meeting,

Padovani, P. 2011, MNRAS, 411, 1547

Perley, R. A., Chandler, C. J., Butler, B. J., & Wrobel, J. M. 2011, ApJ, 739, L1

Piran, T., Nakar, E., & Rosswog, S. 2012, arXiv:1204.6242

Rix, H-W., Barden, M., Beckwith, S. V. W. et al., 2004, ApJS, 152, 163

Sadler, E. M., et al. 2006, MNRAS, 371, 898

Scheers, B. 2011, "Transient and Variable Radio Sources in the LOFAR Sky", PhD thesis, University of Amsterdam.

Schinnerer, E., et al. 2007, ApJS, 172, 46

Seymour, N., et al. 2008, MNRAS, 386, 1695

Smolčić, V., et al. 2008, ApJS, 177, 14

Soderberg, A. M., et al. 2010, Nature, 463, 513

Thyagarajan, N., Helfand, D. J., White, R. L., & Becker, R. H. 2011, ApJ, 742, 49T

Treister, E., Virani, S., Gawiser, E., et al. 2009, ApJ, 693, 1713

Vanzella, E., Cristiani, S., Dickinson, M., et al. 2008, A&A, 478, 83

de Vries, W. H., Becker, R. H., White, R. L., & Helfand, D. J. 2004, AJ, 127, 2565

Welch, J., Backer, D., Blitz, L., et al. 2009, IEEE Proceedings, 97, 1438

Willner, S. P., Ashby, M. L. N., Barmby, P., et al. 2012, ApJ, 756, 72

Wilman, R. J., Jarvis, M. J., Mauch, T., Rawlings, S. Hickey, S., 2010, MNRAS, 405, 447

Wolf, C., Meisenheimer, K., Kleinheinrich, M., et al. 2004, å, 421, 913

Zauderer, B. A., Berger, E., Soderberg, A. M., et al. 2011, Nature, 476, 425

de Zotti, G., Massardi, M., Negrello, M., & Wall, J. 2010, A&A Rev., 18, 1